arXiv:1405.0629v2 [stat.ME] 14 May 2015

Local-Aggregate Modeling for Big-Data via Distributed Optimization: Applications to Neuroimaging

Yue Hu*

Department of Statistics, Rice University

*email: yue.hu@rice.edu

and

Genevera I. Allen*

Dobelman Family Junior Chair

Departments of Statistics and Electrical & Computer Engineering, Rice University,

Department of Pediatrics-Neurology, Baylor College of Medicine,

Jan and Dan Duncan Neurological Research Institute, Texas Children's Hospital.

*email: gallen@rice.edu

Summary: Technological advances have led to a proliferation of structured big data that have matrix-valued covariates. We are specifically motivated to build predictive models for multi-subject neuroimaging data based on each subject's brain imaging scans. This is an ultra-high-dimensional problem that consists of a matrix of covariates (brain locations by time points) for each subject; few methods currently exist to fit supervised models directly to this tensor data. We propose a novel modeling and algorithmic strategy to apply generalized linear models (GLMs) to this massive tensor data in which one set of variables is associated with locations. Our method begins by fitting GLMs to each location separately, and then builds an ensemble by blending information across locations through regularization with what we term an aggregating penalty. Our so called, Local-Aggregate Model, can be fit in a completely distributed manner over the locations using an Alternating Direction Method of Multipliers (ADMM) strategy, and thus greatly reduces the computational burden. Furthermore, we propose to select the appropriate model through a novel sequence of faster algorithmic solutions that is similar to regularization paths. We will demonstrate both the computational and predictive modeling advantages of our methods via simulations and an EEG classification problem.

Key Words: ADMM; Big data; Ensemble learning; Generalized linear models; Multi-subject neuroimaging; Parallel computing; Regularization paths.

1 Introduction

Predictive modeling of subject-level behavior based on whole-brain scans is an important goal of neuroimaging studies. Technologies such as EEG, MEG, and fMRI produce massive amounts of spatio-temporal data (brain images or recordings measured over time) for each subject that serve as covaraites; but most studies typically consist of only tens to hundreds of subjects from which to build supervised predictive models. This problem is thus ultra-high-dimensional and poses both methodological and computational problems for classical statistical and machine learning techniques. First, in $p \gg n$ settings, classical statistical models are ill-posed and cannot be fit without some form of regularization, dimension reduction, or variable selection. Second, the predictors in this supervised learning problem are not vectors as is typical, but are matrices; in other words, our data can be arranged as a 3D array of subjects by brain locations by time points. Third, the size of this big data (for example, $\approx$4GB per subject for fMRI data) can easily exceed the memory limit of mathematical programming tools or even the capacity of a single computer node. Finally, the covariates for each subject are spatio-temporal, and few existing statistical learning methods directly account for these strong structural dependencies. In this paper, our objective is to develop a novel multivariate modeling framework and fast algorithmic strategy for multi-subject neuroimaging studies that addresses each of the above challenges and leads to better predictive performance with scientifically more interpretable results.

Many statistical machine learning techniques have been proposed for supervised modeling of multi-subject neuroimaging data. First, one could apply standard multivariate machine learning techniques to whole-brain data (Pereira et al., 2009), such as support vector machines (De Martino et al., 2007), which typically take a vector of covariates as inputs. To apply these machine learning methods then, the locations by time series for each subject must be vectorized, thus unraveling the important spatio-temporal data structure. Alternatively, one could apply machine learning techniques (De Martino et al., 2007; Calhoun et al., 2009) after first reducing the dimension of each subject's brain scan using methods such as principal components or independent components analysis. While dimension reduction techniques effectively reduce the computational burden of multi-subject modeling, the resulting predictions are not directly interpretable in the original domain. Recently, there has been some interest in fitting supervised models to a tensor of covariates in the statistical literature. The methods of (Hung and Wang, 2013; Zhou et al., 2013; Zhou and Li, 2014) all enforce some form of rank constraints on matrix or tensor coefficients. While low-rank models are a methodological solution to fitting multi-

subject models in ultra-high-dimensional settings, these methods are computationally intensive; for example, the method of (Zhou and Li, 2014) requires computing repeated singular value decompositions. Furthermore, none of these aforementioned approaches or standard machine learning methods directly take the strong spatio-temporal dependencies observed in neuroimaging data into account. We seek a method that both naturally deals with a tensor of covariates and also respects the spatio-temporal nature of whole-brain images.

The development of our modeling strategy will be primarily motivated by computational concerns. Fitting statistical learning models to ultra-high-dimensional data such as whole-brain multi-subject neuroimaging data is a major computational hurdle requiring huge amounts of memory ($\approx$ 4GB per subject) and long processing times. A possible way around this is to break up the multivariate problem into a series of smaller problems that can be fit independently. One could imagine fitting a separate statistical model to each brain location (e.g. voxel); for each of these "local" models, we are back in the common framework of a vector of covariates, the time series, for each subject. This has close connections with the "massive univariate analysis" such as the random effects general linear model (Friston et al., 1994) that is widely used for finding spatial maps of brain activation in fMRI data. Such massive independent analyses can be computed efficiently by using parallel computing such as with Message Passing Interfaces (MPIs) or Graphic Processing Units (GPUs). While these massive univariate frameworks are appealing because their computation can be distributed, they do not result in a unified whole-brain model to predict subject-level attributes. Furthermore, we expect adjacent brain locations to be highly correlated. Massively univariate methods then lose important information that can be gained by considering all brain locations through a multivariate model.

We seek to develop a method that has the computational advantages of massively univariate methods, but still directly accounts for the spatio-temporal tensor structure of multi-subject neuroimaging studies. To this end, we develop a framework that fits separate models to each brain location, but then combines these set of local models in an ensemble that blends information across nearby local models through regularization.

In this paper, we make several methodological and computational contributions. Methodologically, we introduce a novel modeling framework for (1) tensor-valued data as (2) an ensemble of local predictions that (3) directly incorporates spatial and temporal information through regularization. As previously mentioned, this framework fits separate models to each location like the massive univariate frameworks, but then blends information across locations using regularization. Our overall ensemble is then the aggregate or sum of all the individual local models with

the regularization term serving to both smooth local model coefficients spatially and weight the local models according to their predictive ability. While our so-called *Local-Aggregate* modeling framework (Section 2) is general, in this paper, we specifically focus on predictive modeling via Generalized Linear Models (GLMs). Because of the unique structure of our Local-Aggregate model, we can (4) fit this multivariate model in a fully distributed manner (both in terms of computation and more importantly, memory and data storage) via a simple splitting algorithm, the alternating direction method of multipliers (ADMM), a major computational contribution (Section 3). Additionally, our algorithm leads to (5) a novel strategy to speed the model selection process by using what we term an *algorithm path* as an approximation to computationally involved regularization paths (Section 4). We will show through simulations (Section 5) and a real multi-subject EEG classification problem (Section 6) that our methodological contributions lead to a model with improved predictive performance as well as scientifically more interpretable results while our computational contributions can dramatically reduce the computational burden of fitting ultra-high-dimensional models.

2 Methods: Local-Aggregate Modeling

Before we begin, let us review the notation we will use in this paper. Tensors are denoted by $\boldsymbol{\mathcal{X}}$, matrices by $\mathbf{X}$, vectors by $\mathbf{x}$ and scalars by x. It maybe necessary to vectorize a matrix $\mathbf{X} \in \Re^{n \times p}$ as $\text{vec}(\mathbf{X}) \in \Re^{np}$; or matricize a tensor along a particular mode: for $\boldsymbol{\mathcal{X}}$ of size $n \times p \times q$, $\mathbf{X}_{(1)}$ is of size $n \times pq$. Outer products of two vectors will be denoted by $\mathbf{x} \circ \mathbf{y}$, and Kronecker products by $\mathbf{X} \otimes \mathbf{Y}$. The ℓ_2 norm over groups, $g \in \mathcal{G}$, is given by $\sum_{g \in \mathcal{G}} \| \mathbf{x}_g \|_2$.

Our goal is to develop a modeling framework for predicting subject-level responses based on a three-dimensional tensor array of covariates corresponding to a matrix of predictors for each subject. As we are primarily motivated by modeling multi-subject neuroimaging data such as with EEG and fMRI, we denote our subject-level response as $\mathbf{y} \in \Re^n$ for continuous behavioral outcomes or $\mathbf{y} \in \{0,1\}^n$ for binary disease categories, and our predictors as $\boldsymbol{\mathcal{X}}_{n \times \tau \times L}$ for n independent subjects, τ time points and L brain locations. Our primary considerations in developing a new multi-subject predictive model are to directly account for the tensor structure of the predictors, the ultra-high-dimensionality of the problem, and the spatio-temporal nature of neuroimaging data. We also seek to structure our model so that estimation is computationally scalable for big data.

2.1 Local Models

We assume that the generative model for the pair of responses and tensor covariates, $(\mathbf{y}, \boldsymbol{\mathcal{X}})$, follows a matrix generalized linear model (GLM): $g(\boldsymbol{\mu}) = \alpha + \mathbf{X}_{(1)}^T vec(\mathbf{B})$, where $\boldsymbol{\mu} = \mathbb{E}(\mathbf{y} \mid \boldsymbol{\mathcal{X}})$ is the conditional mean responses, $g()$ is the canonical link function associated with a particular exponential family (e.g. the identity link, $g(\mu) = \mu$ for the Gaussian distribution), $\alpha \in \Re$ is the intercept, and the coefficient matrix $\mathbf{B} \in \Re^{\tau \times L} = [\boldsymbol{\beta}_1 \dots \boldsymbol{\beta}_L]$ is the collection of local coefficient vectors $\boldsymbol{\beta}_l \in \Re^{\tau}$. This multivariate model follows the construction of classical GLMs (McCullagh, 1984) treating the vectorized locations by time points as covariates.

These matrix GLMs are ultra-high-dimensional statistical problems with n in the tens to hundreds of subjects relative to $p = L\tau$ on the order of tens of thousands to millions. Hence, several statisticians have recently proposed regularization methods to fit matrix GLMs by using low-rank constraints, penalties, or low-rank structured coefficient models (Zhou and Li, 2014; Zhou et al., 2013) or by placing two-way penalties on the coefficient matrix (Tian et al., 2012). These matrix GLM methods, however, suffer from both computational and statistical inefficiencies. Even with regularization, the number of parameters to be estimated is huge relative to the number of independent observations and thus we cannot expect to accurately estimate matrix GLM parameters. Second, matrix GLMs cannot be fit in an easily distributable manner. Consider a simple linear model, for example, with squared error loss: $\| \mathbf{y} - \mathbf{X}_{(1)} \, \mathrm{vec}(\mathbf{B}) \|_2^2 = \sum_{i=1}^{n} (y_i - \sum_{t=1}^{\tau} \sum_{l=1}^{L} x_{itl} \beta_{tl})^2$; as the matrix parameters are coupled together in the loss function, it is not conducive to distributed optimization.

Thus, we are motivated to consider an alternative way to model the responses and tensor covariates that will approximate the matrix GLM but also be statistically and computationally more efficient. We consider modeling the response at each location separately through local GLMs: $g(\boldsymbol{\mu}_l) = \alpha_l + \mathbf{X}_l \boldsymbol{\beta}_l, \forall \; l = 1, \dots L$. Here, $\mathbf{X}_l \in \Re^{n \times \tau}$ is data at location l, $\boldsymbol{\mu}_l = \mathbb{E}(\mathbf{y} \mid \mathbf{X}_l)$ is the local conditional mean, and $\alpha_l \in \Re$ is the local intercept. We propose to build an ensemble of these local models by aggregating through regularization, which we will discuss in the next section. Notice first, however, that the coefficients in our collection of local GLMs are clearly only equivalent to the coefficients of the matrix GLM if each location is independent. But with neuroimaging data, we expect brain locations to be strongly spatially correlated. A key modeling assumption, then, is that if we account for the spatial dependencies of the parameters (through regularization, for example, discussed subsequently), then the local models are approximately independent.

If our proposed collection of local models are only an approximation to the true matrix

GLM, then why do we expect these models to work well in practice? We answer this question fully in Section 2.3 after we introduce our optimization framework. Briefly, however, notice that each of our local models has τ parameters, far fewer relative to the matrix GLM. Thus, we expect our statistical efficiency in estimating parameters of the local model to be much better than that of the matrix GLM. Second, fitting our models is conducive to simple distributed optimization. Consider again the example of squared error loss, now for the sum of our local models: $\sum_{l=1}^{L} \| \mathbf{y} - \mathbf{X}_l \boldsymbol{\beta}_l \|_2^2 = \sum_{i=1}^{n} \sum_{l=1}^{L} (y_i - \sum_{t=1}^{\tau} x_{itl} \beta_{tl})^2$; this loss is location-separable and hence easily distributable across locations. Finally, our collection of local models can be viewed as an ensemble learning method, and ensembles are known to yield powerful predictive models.

2.2 Local-Aggregate Optimization Framework

We propose to build an ensemble learning method by aggregating our local models through regularization. Fitting a simple sum of local models, however, does not account for the spatial structure of neuroimaging data. Thus, we seek a method of combining our local models in such a way to borrow strength across neighboring locations through an *aggregating penalty*, $P_{agg}(\mathbf{B}, \mathbf{G})$ with external spatial information matrix $\mathbf{G}$ discussed subsequently. Also in many applications including neuroimaging, each of the local GLMs could be high-dimensional if $\tau > n$, and thus, some local regularization may also be needed. Then, our general Local-Aggregate optimization framework below is an ensemble of regularized local GLMs with an aggregating penalty:

$$\underset{\boldsymbol{\alpha}, \mathbf{B}}{\text{minimize}} \quad \underbrace{\sum_{l=1}^{L} [\underbrace{\ell(\mathbf{y}; \alpha_l + \mathbf{X}_l \boldsymbol{\beta}_l)}_{\text{GLM loss}} + \underbrace{\lambda_{loc} P_{loc}(\boldsymbol{\beta}_l)}_{\text{Possible local penalties}}]}_{\text{Sum of penalized log-likelihoods for local GLMs}} \underbrace{+\lambda_{agg} P_{agg}(\mathbf{B}, \mathbf{G})}_{\text{Aggregating over locations}}. \tag{2.1}$$

Here, $\ell()$ is the GLM loss function, or the negative log-likelihood, $P_{loc}(\boldsymbol{\beta}_l) : \Re^{\tau} \to \Re^{+}$ is a local penalty, $P_{agg}(\mathbf{B}, \mathbf{G}) : (\Re^{\tau \times L}, \Re^{\tau \times \tau}) \to \Re^{+}$ is the aggregating penalty, and $\lambda_{loc} \geq 0, \lambda_{agg} \geq 0$ are tuning parameters of local and aggregating regularization respectively.

2.2.1 Aggregating Penalty

To account for the spatial dependencies in the data, we propose to enforce the local coefficients to be smooth over locations through our aggregating penalty. Popular approaches to achieve spatial smoothness include imposing smoothness with respect to a graph structure through the

graph penalty, $|| \boldsymbol{\beta} ||^2_{\mathbf{G}} = \boldsymbol{\beta}^T \mathbf{G} \boldsymbol{\beta}$ (Grosenick et al., 2013; Allen et al., 2014). Taking $\mathbf{G}$ to be the graph Laplacian (the difference between the degree matrix $\mathbf{D} \in \Re^{L \times L}$ and the adjacency matrix $\mathbf{W} \in \Re^{L \times L}$, where $\mathbf{W}$ is a symmetric matrix with nonnegative entries $w_{l,l'}$ and $\mathbf{D}_{l,l'} = \sum_{l=1}^{L} w_{l,l'}$), the graph penalty has the appealing interpretation of regularizing pairwise differences between coefficients that are adjacent in the graph: $|| \boldsymbol{\beta} ||^2_{\mathbf{G}} = \sum_{(i,j) \in \mathbf{G}} (\beta_i - \beta_j)^2$ (Grosenick et al., 2013). Consider a multivariate extension of a weighted graph Laplacian penalty where we penalize the weighted differences between coefficient vectors: $\sum_{l \neq l'} w_{l,l'} || \boldsymbol{\beta}_{\cdot l} - \boldsymbol{\beta}_{\cdot l'} ||^2_2$, where $\boldsymbol{\beta}_{\cdot l} \in \Re^{\tau}$ is the l^{th} column of $\mathbf{B}$, and $w_{l,l'}$'s are weights inversely proportional to the distance between locations l and l' (in subsequent sections, we will use $\boldsymbol{\beta}_l$ to denote $\boldsymbol{\beta}_{\cdot l}$ for notational convenience). Note that if we take $\mathbf{G}$ to the be the $L \times L$ weighted graph Laplacian, $\mathbf{G} = \mathbf{D} - \mathbf{W}$, then the multivariate graph Laplacian penalty is equal to $tr(\mathbf{B}\,\mathbf{G}\,\mathbf{B}^T)$, i.e., $tr(\mathbf{B}\,\mathbf{G}\,\mathbf{B}^T) = \sum_{(l \neq l')} w_{l,l'} || \boldsymbol{\beta}_{\cdot l} - \boldsymbol{\beta}_{\cdot l'} ||^2_2$.

Hence, we can use an aggregating penalty of the form $P_{agg}(\mathbf{B}, \mathbf{G}) = \text{tr}(\mathbf{B}\,\mathbf{G}\,\mathbf{B}^T)$, noting that we can specify $\mathbf{G}$, and more specifically the weights $w_{l,l'}$ defining $\mathbf{G}$, in such a way as to interpret our penalty as a graph type penalty. Moreover, as $\sum_{(l \neq l')} w_{l,l'} || \boldsymbol{\beta}_{\cdot l} - \boldsymbol{\beta}_{\cdot l'} ||^2_2 = \sum_{t=1}^{\tau} \boldsymbol{\beta}_{t\cdot}^T \mathbf{G}\, \boldsymbol{\beta}_{t\cdot}$, we can also interpret $\mathbf{G}$ as a roughness penalty matrix often used for smoothing in functional data analysis (Ramsay and Silverman, 2005).

We are then left with the task of choosing the weights $w_{l,l'}$, which is application specific, but there are two general classes of well studied and employed weighting schemes:

- $w_{l,l'} = \begin{cases} 1 & \text{if } l' \in \mathcal{N}(l) \\ 0 & \text{otherwise} \end{cases}$, where $\mathcal{N}(l)$ is the neighborhood of node l. This has been employed when locations are laid out on a regular grid as in fMRI data (Grosenick et al., 2013), or when locations correspond to an existing network structure (Huang et al., 2009).
- $w_{l,l'} = Kern(D_{l,l'}, \theta)$, where $Kern()$ is a kernel smoothing function that takes $D_{l,l'}$, the distance between locations l and l'. While there are many examples of appropriate kernel smoothers (Ramsay and Silverman, 2005), a common smoother is the exponential kernel, $w_{l,l'} = e^{-D^2_{l,l'}/\theta}$. These types of weights are useful for data where the locations correspond to some irregular physical coordinates, such as channels of EEG or MEG signals.

2.2.2 Example: Local-Aggregate Modeling for Multi-Subject Neuroimaging Data

We pause to discuss some possible forms our Local-Aggregate modeling framework may take for our primary motivating example of multi-subject spatio-temporal neuroimaging data such as with EEG, MEG, or fMRI data. As measurements are taken over time in these examples, we also expect the coefficients to be smooth with respect to time. Therefore, we may wish to encourage

temporal smoothness through a local quadratic smoothing penalty $P_{loc}^{sm}(\boldsymbol{\beta}_l) = \boldsymbol{\beta}_l^T \boldsymbol{\Omega} \boldsymbol{\beta}_l$. where $\boldsymbol{\Omega}$ is the second order difference matrix penalizing the squared difference between coefficients at adjacent time points (Eilers and Marx, 1996). Additionally, some neuroimaging data consists of many locations; consider EEG data which can have up to 512 channels, not all of which are expected to contribute to subject-level responses. Hence, we may also want to encourage sparsity to select only the relevant locations for prediction purposes. However, since our local coefficient, $\boldsymbol{\beta}_l$, is a vector, we cannot simply apply a scalar penalty, such as the ℓ_1 norm. Thus, we suggest a group lasso penalty, $\sum_{l=1}^{L} \| \boldsymbol{\beta}_l \|_2$ (Yuan and Lin, 2006), which treats coefficients at each location as a group, either zeroing out all elements of $\boldsymbol{\beta}_l$ or letting all elements be non-zero. Also, this group lasso penalty can easily be localized, $P_{loc}^{sp}(\boldsymbol{\beta}_l) = \| \boldsymbol{\beta}_l \|_2$, as it is separable in locations. Putting everything together, we have the Local-Aggregate optimization framework for neuroimaging data:

$$\underset{\boldsymbol{\alpha},\mathbf{B}}{\text{minimize}} \quad \sum_{l=1}^{L} \left[\ell(\mathbf{y}; \alpha_l + \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc}^{sm} \boldsymbol{\beta}_l^T \boldsymbol{\Omega} \boldsymbol{\beta}_l + \lambda_{loc}^{sp} \| \boldsymbol{\beta}_l \|_2 \right] + \lambda_{agg} \text{tr}(\mathbf{B}\,\mathbf{G}\,\mathbf{B}^T), \tag{2.2}$$

where $\lambda_{loc}^{sm} \geq 0$ and $\lambda_{loc}^{sp} \geq 0$ are local tuning parameters, and the choice of $\mathbf{G}$ depends on the spatial information of the specific imaging modality as discussed in Section 2.2.1.

2.2.3 Subject-level Behavior Prediction

Once we obtain the estimates of all parameters, we propose to predict the subject-level response by taking the average of the local predictions: $\hat{\mathbf{y}} = \frac{1}{L} \sum_{l=1}^{L} g^{-1}(\hat{\alpha}_l + \mathbf{X}_l \hat{\boldsymbol{\beta}}_l)$, where $g^{-1}()$ is the inverse link function associated with the exponential family distribution. Thus, our overall prediction is an ensemble of local predictions. Note if our local models were estimated completely separately as with $\lambda_{agg} = 0$, then this ensemble would be the naive raw average or majority vote of local models. On the other hand, by using our aggregating penalty and estimating the local models together, the local coefficients that determine the local predictions depend on each other; hence, the coefficients of locations less useful for prediction are shrunk towards zero. Overall, the predictions of our Local-Aggregate model behave like an ensemble.

2.3 Advantages of Local-Aggregate Modeling Framework

Our method offers several advantages in terms of statistical efficiency (**S**), modeling flexibility (**M**), and computational efficiency (**C**). Here, we discuss some of these advantages as to existing

literature.

(S.1) Predictive power of ensembles. The idea of ensemble learning is simple: combine strengths of multiple base models or weaker learners to build a set of models with superior predictive accuracy. Examples of such learning techniques include Bayesian model averaging, model stacking, bagging, and boosting (Hastie et al., 2009). Most of these existing ensemble learning methods use sampling or iterative reweighting schemes to create a diverse set of base learners for the same data set. Our ensemble of local models, on the other hand, uses separate local data, $\mathbf{X}_l$, for each base model and hence there is no need for computationally-intensive sampling schemes. Meanwhile, our ensemble incorporates external spatial information through spatially smoothed $\hat{\boldsymbol{\beta}}_l$, resulting in a combination of weaker yet more informative local learners that can potentially yield better prediction than one sophisticated multivariate model.

(S.2) Lower-dimensional problem. The matrix GLM problem is ultra-high-dimensional with small n relative to $L\tau$ parameters as compared to τ parameters for each of our local models. Thus, we can achieve much improved statistical efficiency to estimate the parameters of our local models than for the true matrix GLM. Often obtaining better estimates of incorrect models leads to better predictive accuracy than poorer estimates of correct models. For example, it is well known that biased regularized linear models have better prediction accuracy under much weaker conditions than the correct, unbiased linear model (Hoerl and Kennard, 1970). Since our Local-Aggregate model seamlessly translates an ultra-high-dimensional problem to a lower-dimensional one, we expect that the improved statistical efficiency will result in better predictive performance. Our method is also highly scalable: as the number of locations L increases, the complexity of our local models remains constant.

(M.1) Spatio-temporal smoothness. Our Local-Aggregate modeling framework directly accounts for the spatio-temporal structure of the data by encouraging both spatial and temporal coefficient smoothness, thus yielding more interpretable results. This is similar to the two-way smoothing penalizes used in the context of dimension reduction by Huang et al. (2009); Allen et al. (2014) and MEG signal reconstruction by Tian et al. (2012).

(M.2) Location selection. Our modeling framework can select brain locations important for prediction in a data-driven manner through the group lasso penalty; this shrinks the coefficients at non-informative locations towards zero and their local predictions to $\hat{\alpha}_l$. Therefore the aggregating penalty combined with the group lasso penalty automatically weight the local coefficient vectors by their predictive power, and our overall prediction is unaffected by summing over uninformative locations. When considering the advantages of our method, one may question the

relative roles of modeling flexibility as compared to statistical efficiency. To address this directly, we study our ensemble of local models and the full matrix GLM with the same local and aggregating penalties in the Supplementary Material Section C. Briefly, our investigation suggests that the improvements in predictive performance we observe are due to both the statistical and modeling advantages of our approach.

(C) Parallelizable Algorithm. Distributed storage and computation is crucial to big data modeling, especially when the data is too large to fit into one computer's memory. However, few statistical algorithms are able to distribute data storage and memory at the same time. Our Local-Aggregate modeling framework is intrinsically primed for parallel computing, discussed subsequently, as the only non-separable term is the aggregating penalty.

3 Algorithm

Our primary goal is to build a scalable modeling and algorithmic strategy for structured big data with many locations, more specifically, large-scale regression problems with matrix-valued predictors. Notice that as our objective function (2.1) is jointly convex in $\mathbf{B}$ and $\boldsymbol{\alpha}$, there are many potential optimization strategies that enable parallel computing; however, many of these do not allow for distributed memory and data storage. We desire a highly-parallelizable algorithm that permits distributed data storage and memory with minimal message passing between the computing nodes. To achieve this, we turn to the Alternating Direction Method of Multipliers (ADMM) optimization strategy.

3.1 Local-Aggregate ADMM

To develop our Local-Aggregate algorithm, we begin with no local intercept α_l to match the classical ADMM framework outlined in Boyd et al. (2011):

$$\underset{\mathbf{B}}{\text{minimize}} \quad \sum_{l=1}^{L} \left[\ell(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l)\right] + \lambda_{agg} P_{agg}(\mathbf{B}, \mathbf{G}). \tag{3.1}$$

We will show how to deal with local intercepts in detail later in Section 3.2.

In order to apply ADMM to our Local-Aggregate modeling framework, we first create a copy $\mathbf{Z} \in \Re^{\tau \times L}$ to substitute the modeling coefficient matrix $\mathbf{B}$ in the aggregating penalty so that we can split the original problem (3.1) into the local models and the aggregating term. Then we

require an additional constraint $\mathbf{B} = \mathbf{Z}$ to ensure the two copies are the same:

$$\underset{\mathbf{B},\mathbf{Z}}{\text{minimize}} \quad \sum_{l=1}^{L} \left[\ell(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l)\right] + \lambda_{agg} P_{agg}(\mathbf{Z}, \mathbf{G}) \quad \text{subject to } \mathbf{B} = \mathbf{Z} \tag{3.2}$$

The subproblems are solved by minimizing the *augmented Lagrangian* of (3.2) with respect to $\mathbf{B}$ and its copy $\mathbf{Z}$ independently:

$$\mathcal{L}_\rho(\mathbf{B}, \mathbf{Z}, \mathbf{U}) = \sum_{l=1}^{L} \left[\ell(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l)\right] + \lambda_{agg} P_{agg}(\mathbf{Z}, \mathbf{G}) + \sum_{l=1}^{L} \frac{\rho}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l + \mathbf{u}_l \|_2^2, \tag{3.3}$$

where $\mathbf{U} \in \Re^{\tau \times L}$ is the dual variable, and $\rho > 0$ is an algorithm parameter. The solutions to the subproblems are then coordinated via a dual update procedure to find the global solution. Thus, the three key steps of our Local-Aggregate ADMM are as follows:

1. $\boldsymbol{\beta}_l^{k+1} = \arg\min_{\boldsymbol{\beta}_l} \ell(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l) + \frac{\rho}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l^k + \mathbf{u}_l^k \|_2^2 \rightarrow \mathbf{B}$-subproblem (in parallel)
2. $\mathbf{Z}^{k+1} = \arg\min_{\mathbf{Z}} \lambda_{agg} P_{agg}(\mathbf{Z}, \mathbf{G}) + \frac{\rho}{2} \sum_{l=1}^{L} \| \mathbf{z}_l - \boldsymbol{\beta}_l^{k+1} - \mathbf{u}_l^k \|_2^2 \rightarrow \mathbf{Z}$-subproblem (message passing)
3. $\mathbf{u}_l^{k+1} = \mathbf{u}_l^k + \boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1} \rightarrow$ dual update (in parallel)

We are fitting regularized GLMs to each location, for which many well-studied optimization algorithms exist (Friedman et al., 2010). If the local penalty is smooth, such as with a temporal roughness penalty, we can apply a fast Newton algorithm to solve the $\mathbf{B}$-subproblem, which achieves at least quadratic convergence rate. If the local penalty is non-smooth, we then employ the proximal gradient method, which guarantees at least $\mathcal{O}(1/k)$ convergence rate under the assumption of convex loss functions (at least one being strongly convex) (Nesterov et al., 2007). Once the local regression information from previous $\mathbf{B}$ updates is gathered, the aggregating penalty enforces smoothness of coefficients over neighboring locations. The $\mathbf{Z}$-subproblem collects the local coefficients and blends structural information across nearby locations through the spatial smoothness penalty, $P_{agg}(\mathbf{B}, \mathbf{G})$. The smoother structural information is then distributed to all locations for the next dual update step. The dual variable $\mathbf{U}$ tracks the progress of the algorithm as $\mathbf{B}$ is squeezed closer to $\mathbf{Z}$ until the equality constraint is satisfied. The dual update step plays an important role of coordinating $\mathbf{B}$ and $\mathbf{Z}$ towards the global solution: as the algorithm converges, $\mathbf{Z}$ approaches $\mathbf{B}$.

Note that as the augmented Lagrangian (3.3) is separable in locations, we can solve the ensemble of separate local models ($\boldsymbol{\beta}_l$-subproblems) completely in parallel. Another major computational advantage of our algorithm is its distributed data storage and minimal message

passing among computing nodes. To start with, we send the response $\mathbf{y}$ and the local covariates $\mathbf{X}_l$ to each computing node l to compute the $\boldsymbol{\beta}_l$ updates, which will be collected to solve the $\mathbf{Z}$ subproblem, then we only need to pass back the new $\mathbf{z}_l$'s, as all other variables are local. Finally, the dual variable update step is local and requires no message passing.

3.1.1 Convergence Analaysis

Our Local-Aggregate ADMM is guaranteed to converge:

Theorem 3.1. *(Extension of a result from Section 3.2 in Boyd et al. (2011)) Under the following assumptions,* (a) $\ell() : \Re^{\tau \times L} \to \Re$, $P_{loc}() : \Re^{\tau} \to \Re^{+}$ *and* $P_{agg}() : \Re^{\tau \times L} \to \Re \cup \{+\infty\}$ *are closed, proper and convex,* (b) $\rho, \lambda_{loc}, \lambda_{agg} \in \Re^{+}$ *fixed, the ADMM iterates satisfy the following convergence: 1) Residual convergence:* $\mathbf{B}^k - \mathbf{Z}^k \to 0$ *as* $k \to \infty$. *2) Objective convergence:* $\sum_{l=1}^{L}[\ell(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l^k) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l^k)] + \lambda_{agg} P_{agg}(\mathbf{Z}^k, \mathbf{G}) \to p^*$, *the optimal value. 3) Dual convergence:* $\mathbf{U}^k \to \mathbf{U}^*$ *as* $k \to \infty$, *where* $\mathbf{U}^*$ *is a dual optimal point scaled by* ρ.

The convergence of the algorithm hinges upon the convexity of the GLM loss function and the penalty functions. As all GLMs with canonical link functions are convex, and many commonly used penalty functions, e.g. lasso, ridge, elastic net, group lasso, are all convex functions, we have a very flexible choice of local GLMs and local penalty functions, and hence our Local-Aggregate ADMM algorithm is suitable for many applied statistical problems.

The convergence rate of our Local-Aggregate ADMM is $\mathcal{O}(1/k)$, which is the same as that of ADMM (Deng and Yin, 2012), which can be further accelerated by using a variable penalty parameter ρ depending on the algorithm progress (how fast $\mathbf{B}$ is pushed towards $\mathbf{Z}$) (Boyd et al., 2011). We also develop a self-adaptive vectorized penalty parameter updating scheme which takes into account the difference in convergence speeds among all time points and automatically selects appropriate ρ for each time point. This further accelerates the algorithm and is described in detail in the supplementary material.

3.2 Local-Aggregate ADMM Example: Multi-Subject Neuroimaging Data

To illustrate the details of our algorithm framework, we study our neuroimaging inspired problem (2.2), for binary classification with local logistic loss functions. (This is also the model we employ for an EEG classification example in Section 6). Let $\mathbf{x}_{il} = \boldsymbol{\mathcal{X}}_{(i,:,l)} \in \Re^{\tau}$ denote measurements for subject i taken at location l, and $\mathbf{y} \in \Re^{n}$ denote subject-level responses, we then have the

following neuroimaging classification problem:

$$\underset{\boldsymbol{\alpha},\mathbf{B},\mathbf{Z}}{\text{minimize}} \sum_{l=1}^{L}[\sum_{i=1}^{n}\{-\mathbf{y}_i(\alpha_l + \mathbf{x}_{il}^T\boldsymbol{\beta}_l) + \log(1+e^{(\alpha_l+\mathbf{x}_{il}^T\boldsymbol{\beta}_l)})\} + \lambda_{loc}^{sm}\,\boldsymbol{\beta}_l^T\,\boldsymbol{\Omega}\,\boldsymbol{\beta}_l + \lambda_{loc}^{sp}\|\,\boldsymbol{\beta}_l\,\|_2] + \lambda_{agg}\,\mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T), \quad \text{subject to } \mathbf{B}=\mathbf{Z}\,. \tag{3.4}$$

However, as the classical ADMM algorithm only allows for two coupled sets of variables (Chen et al., 2013), we need to transform the above problem into the Local-Aggregate ADMM framework with no local intercepts through reparametrization. Let $\widetilde{\mathbf{x}}_{il} = [1;\mathbf{x}_{il}] \in \Re^{(\tau+1)}$, $\widetilde{\boldsymbol{\beta}}_l = [\alpha_l;\boldsymbol{\beta}_l] \in \Re^{(\tau+1)}$, $\mathbf{S} = \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_\tau \end{pmatrix}$, $\widetilde{\boldsymbol{\Omega}} = \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Omega} \end{pmatrix}$, both $\Re^{(\tau+1)\times(\tau+1)}$ matrices. Let $\mathrm{Prox}_g(x,t) = I_{\{\|x\|_2 \geq t\}}(1-t/\|x\|_2)x$ denote the proximal operator of ℓ_2 norm (Yuan and Lin, 2006), let $f(\boldsymbol{\beta}_l)$ denote the logistic loss plus the the local smoothness penalty, and let $\mathbf{r} = \boldsymbol{\beta}_l - \mathbf{z}_l$ denote the primal residual, and $\mathbf{s} = \mathbf{z}^k - \mathbf{z}^{k-1}$ denote the dual residual. Then, ignoring the $\widetilde{}$ sign for notational convenience, the Local-Aggregate ADMM for solving this classification problem (B.1) is outlined by Algorithm 1. (We formally prove the equivalence of this reparametrized framework to the original problem and its convergence rate in the supplementary material.)

Algorithm 1 : Local-Aggregate ADMM for Neuroimaging Example

Initialize $\mathbf{B}^0 = \mathbf{0}$, $\mathbf{U}^0 = \mathbf{0}, \mathbf{Z}^0 = \mathbf{0}, \gamma \in (0,1)$, $k=1$, ρ initialized.
While $\mathbf{r} > \epsilon^{tol}$ and $\mathbf{s} > \epsilon^{tol}$
  For $l = 1:L$ → Parallel
    While $\|\,\boldsymbol{\beta}_l^k - \boldsymbol{\beta}_l^{k-1}\,\| > \epsilon^{tol}$
      $t = 1$
      $\nabla f^k = \mathbf{X}_l^T(\mathbf{y} - 1/e^{-\mathbf{X}_l\boldsymbol{\beta}_l^k}) + 2\lambda_{loc}^{sm}\,\boldsymbol{\Omega}\,\boldsymbol{\beta}_l^k + \rho\,\mathbf{S}(\boldsymbol{\beta}_l^k - \mathbf{z}_l^k + \mathbf{u}_l^k)$
      $\boldsymbol{\beta}_l^{k+1} = \mathrm{Prox}_g(\boldsymbol{\beta}_l^k - t\nabla f^k, \lambda_{loc}^{sp}t)$
      While $f(\boldsymbol{\beta}_l^k - t\nabla f^k) > f(\boldsymbol{\beta}_l^k) - t(\nabla f^k)^T\mathrm{Prox}_g(\boldsymbol{\beta}_l^k - t\nabla f^k, \lambda_{loc}^{sp}t) +$
        $\frac{t}{2}\|\mathrm{Prox}_g(\boldsymbol{\beta}_l^k - t\nabla f^k, \lambda_{loc}^{sp}t)\|_2^2$
          $t := \gamma t$
          $\boldsymbol{\beta}_l^{k+1} = \mathrm{Prox}_g(\boldsymbol{\beta}_l^k - t\nabla f^k, \lambda_{loc}^{sp}t)$
      End while
    End while
  End for → End parallel
  $\mathbf{Z}^{k+1} = \rho(\mathbf{S}\,\mathbf{B}^{k+1} + \mathbf{U}^k)(\rho\,\mathbf{I}_L + 2\lambda_{agg}G)^{-1}$ → Message passing
  For $l = 1:L$ → Parallel
    $\mathbf{u}_l^{k+1} = \mathbf{u}_l^k + \boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1}$
  End for → End parallel
  $k = k+1$
End while

Overall, we have developed a highly-parallelizable ADMM algorithm with distributed memory and data storage, minimal message passing, and fast convergence speed, all of which are well suited for our Local-Aggregate optimization framework for multi-subject neuroimaging data with many locations.

4 Model Selection: Local-Aggregate Algorithm Path

We have carefully developed a fast, completely parallelizable algorithmic strategy for large-scale matrix-covariate regression problems. In practice, however, often the most computationally intensive part of fitting statistical machine learning models is not fitting at a single regularization parameter, λ, but instead performing model selection by considering a sequence of λ's. Typically, one considers a grid of $\boldsymbol{\lambda} = \{\lambda = 0, \ldots, \lambda_{max}\}$ and chooses a value of λ that minimizes a selection criterion such as BIC, AIC, GCV, or employs sampling schemes such as cross-validation (CV) or stability selection at each λ. For our problem, repeatedly fitting 50-100 models (or more for re-sampling) is computationally infeasible. Thus, we consider a completely novel yet effective approach to model selection.

First, for our Local-Aggregate modeling framework, notice that we can select any λ_{loc} completely separately for each location in parallel. Thus we are left with how to choose the aggregating penalty, λ_{agg}, which cannot be done in a parallel manner. Our goal is to develop a computationally efficient method to select the correct amount of smoothing over locations. Let us start with a simulated linear regression example: suppose we have a chain graph of $L = 10$ locations, $n = 100$ subjects, and $\tau = 20$ time points. The classical regularization path is computed by varying λ_{agg}, and the entire path of solutions is shown at $L = 1, 5$, and 10 in top panel of Figure 1. The regularization path of our Local-Aggregate modeling framework $\{\hat{\mathbf{B}}(\lambda_{agg}) : 0 < \lambda_{agg} < \infty\}$ starts from $\lambda_{agg} = 0$, i.e. completely independent local GLMs, to $\lambda_{agg} = \lambda_{max}$, some large value giving the extreme smoothness over locations. The parameter λ_{agg} controls the amount of smoothing over locations: as λ_{agg} increases, there is increased smoothness over locations.

If we take a closer look at the iterates of our Local-Aggregate ADMM algorithm when we set $\lambda_{agg} = \lambda_{max}$ plotted in the bottom panel of Figure 1, we see that the iterates of our algorithm look very similar to the regularization path for different λ values. In particular, for almost all $\hat{\mathbf{B}}(k)$ (our estimate from the $\mathbf{B}$-subproblem at iterate k of the ADMM algorithm), there exits some corresponding $\hat{\mathbf{B}}(\lambda_j)$ in the regularization path that precisely matches $\hat{\mathbf{B}}(k)$. This observation

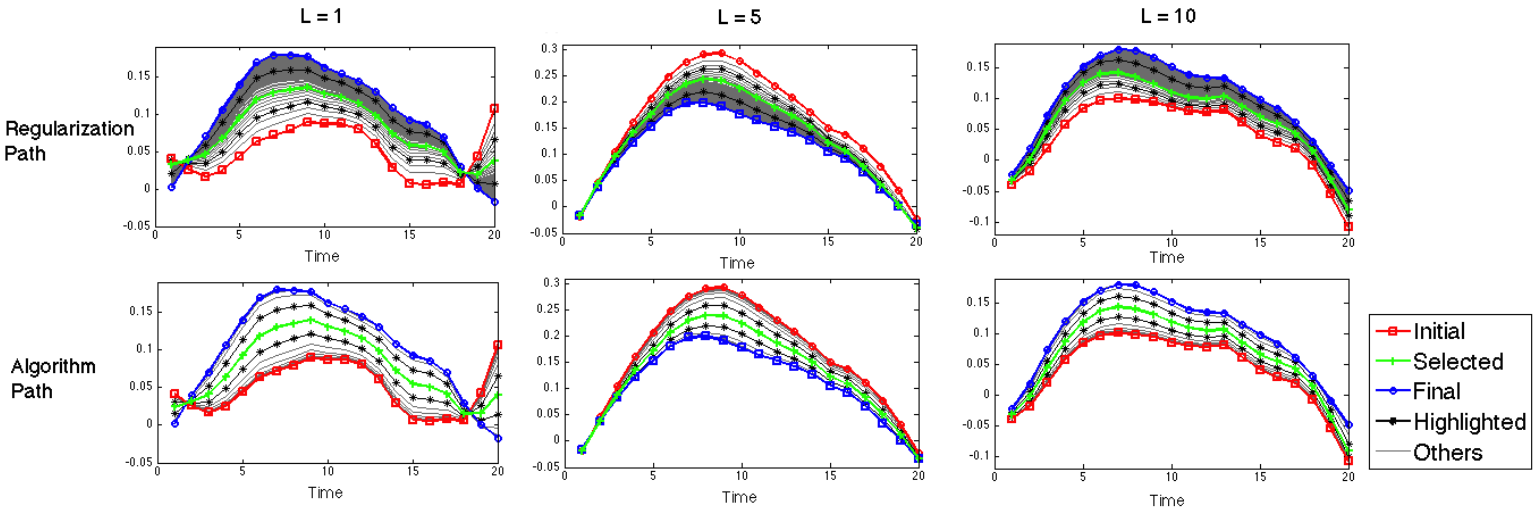

Figure 1: Comparison of the entire regularization path with the algorithm path at L = 1 (left panels), L = 5 (middle panels), and L = 10 (right panels). "Initial" refers to the model at $\lambda = 0$ or $k = 1$, "Selected" refers to the optimal model selected via cross validation, "Final" refers to the model at $\lambda = \lambda_{max}$ or $k = K$, and "Highlighted" refers to some highlighted model along the path for illustrative purposes, and "Others" refers to the rest of the path of solutions. This empirically demonstrates that our algorithm path well approximates the regularization path. Also the model selected by CV is comparable for both methods.

makes sense considering how our Local-Aggregate ADMM algorithm progresses. Namely, at iteration 1, we start with $\mathbf{Z} = \mathbf{0}$, giving no smoothness and a fully local solution for $\mathbf{B}$. Thus at iteration 1, $\hat{\mathbf{B}}(k)$ is equal to the regularization path for $\hat{\mathbf{B}}(\lambda_j = 0)$. As the ADMM algorithm progresses, $\hat{\mathbf{B}}(k)$ gets squeezed towards the overly smoothed $\mathbf{Z}$, meaning that at each subsequent iteration, more smoothness is induced. This occurs until finally the ADMM algorithm converges to the solution at $\lambda_{agg} = \lambda_{max}$, meaning that $\hat{\mathbf{B}}(k = K^{\text{final}}) = \hat{\mathbf{B}}(\lambda_j = \lambda_{max})$. While we have observed this correspondence between our ADMM iterates and the regularization path over a grid of λ's in the simulation in Figure 1 as well as all other empirical investigations, proving such equivalence is beyond the score of this paper; we hope to investigate this formally in future work.

Hence based on our empirical observations, we propose the so-called Local-Aggregate Algorithm Path, which approximates the solutions to the regularization path over a range of λ's with the iterates of our Local-Aggregate ADMM algorithm with fixed tuning parameter at $\lambda_{agg} = \lambda_{max}$. The difference between the regularization path and the local-aggregate algorithm path is illustrated as follows:

• **Regularization Path** Given a sequence of regularization parameters $\lambda_N = \lambda_{max} > \cdots > \lambda_j > \cdots > \lambda_1 = 0$, we solve a sequence of minimization problems corresponding to each λ_k and obtain $\{\hat{\mathbf{B}}(\lambda_j) : j = 1, \cdots, N\}$.

• **Algorithm Path** Given the fixed regularization parameter $\lambda_{agg} = \lambda_{max}$, we solve one minimization problem and take the series of algorithm iterates as our solution path: $\{\hat{\mathbf{B}}(k) : k = 1, \cdots, K\}$, where K is the number of iterations.

While the grid search method evaluates the optimization problem over a sequence of λ_j's, our local-aggregate algorithm path approximates the solutions as the algorithm iterates without tuning the regularization parameter λ_{agg}. The computational cost is thus greatly reduced: namely, from solving N optimization problems to solving just one problem.

4.1 Model Selection via Local-Aggregate Algorithm Path

The computational cost is huge if we employ M-fold cross-validation along with the regularization path to select the optimal model: we need to solve a sequence of N minimization problems for all M cross-validated folds, totaling $M \times N$ model fits. However, if we combine the idea of the Local-Aggregate algorithm path with cross-validation (Algorithm 2), the number of model fits is cut down to only M. Unlike traditional model selection methods, which seek to choose a λ^*, our

Algorithm 2 Compute Local-Aggregate Algorithm Path via Cross Validation

1. Set λ_{agg} large, split data into M folds
2. For $m = 1, \cdots, M$, run local-aggregate algorithm on subsets of data with the m^{th} fold left out. Record $\hat{\mathbf{B}}(k)$ for each iteration of the algorithm.
3. Compute CV error for $\hat{\mathbf{B}}(k)$ generated by each iterate on the m^{th} left out fold.
4. Average the CV errors over M folds, and obtain the CV error curve across iterates of the algorithm.
5. Pick the optimal iteration K^{opt} that has the minimum CV error.

proposed model selection strategy chooses the correct model from one of the algorithm iterates. While we have not formally proven this, we conjecture that our algorithm iterates correspond to the solution at a specific λ value, and thus our model selection procedure is equivalent to M-fold CV. Overall, our algorithm path and model selection procedure represent a completely novel approach that offer substantial computational savings.

5 Simulation

To better understand the performance of our Local-Aggregate modeling framework, and how it compares with the classical methods, we present a series of simulated examples inspired by neuroimaging data. The true spatio-temporal signal $\mathbf{B}_o = \sum_{r=1}^{R} \mathbf{v}_r \, \mathbf{u}_r^T$ is simulated by the sum of outer products of a spatial factor, $\mathbf{u}_r \in \Re^L$ corresponding to a $\sqrt{L} \times \sqrt{L}$ grid consisting of zeros except for two blocks (one near either end of the grid's diagonal), and a temporal factor $\mathbf{v}_r \in \Re^{200}$ with 200 equally-spaced time points following sinusoidal curves, for $r = 1, \ldots, R$, where $R = \text{rank}(\mathbf{B}_o)$. Note that $\{\mathbf{v}_1, \cdots, \mathbf{v}_R\}$ and $\{\mathbf{u}_1, \cdots, \mathbf{u}_R\}$ are sets of orthogonal vectors so that the rank of the true signal is R. The tensor-valued covariates $\boldsymbol{\mathcal{X}}$, are a collection of covariates $\mathbf{X}_i \overset{iid}{\sim} N(0, \Sigma_L \otimes \Sigma_T)$ for each subject, i, generated as matrix-variate normal with exponential spatial ($\Sigma_L = e^{-\mathbf{D}_L^2/\theta_L}$) and temporal ($\Sigma_T = e^{-\mathbf{D}_T^2/\theta_T}$) covariance structures. Here, $\mathbf{D}_L$ and $\mathbf{D}_T$ are spatial and temporal distances respectively, and θ_L and θ_T control the amount of spatial and temporal correlation.

5.1 Regression simulation

For our regression simulation, the response variable is generated as $\mathbf{y} = \mathbf{X}_{(1)}\, vec(\mathbf{B}_o) + \epsilon$, for $\epsilon \overset{iid}{\sim} N(0,1)$. We compare our method to: the Lasso, Ridge, Elastic Net, and nuclear-norm regularized tensor regression method (Zhou et al., 2013). Since the classical regression methods take vectors as covariates, the time series by locations are vectorized. We demonstrate our findings in three examples. First, we simulate data with differing numbers of locations, keeping the rank of the spatio-temporal signal the same. Second, we compare methods for different ranks of $\mathbf{B}_o$ with fixed numbers of locations, $L = 100$. Third, we explore the effects of the amount of spatio-temporal correlation in the covariates for the case of $L = 100$, and $\text{rank}(\mathbf{B}_o) = 2$.

We present simulation results for fifty replicates and compare various methods in terms of both prediction accuracy and signal recovery. We use mean squared error (MSE) as the criterion for prediction error, and to evaluate signal recovery, we use MSE of the estimated coefficient matrix, $\|\hat{\mathbf{B}} - \mathbf{B}_o\|_2$, and the true/false positive rates of the true location detection.

5.1.1 Number of Locations. Our first set of simulations varies the number of locations: $L = 25, 64, 100, 144$. The results shown in Table 1 (top) demonstrate that our Local-Aggregate modeling framework outperforms other methods when the number of locations is large. Note that as $p = L \times \tau = 28800 \gg n = 100$, this represents an ultra-high-dimensional problem, and the performance of other methods in terms of prediction error declines as L increases. The better performance of our method when L is large can be easily explained: as the number of base models increases, our ensemble is able to predict the subject-level response more accurately. Additionally, when L is large, more external spatial information is brought into the modeling framework, therefore we expect more accurate predictive results.

5.1.2 Signal Complexity. The signal complexity is simulated as the ranks of the original signal $\mathbf{B}_o$. In the rank 2 case, the spatial signal has two major areas of interest, and the corresponding temporal signals are generated by 200 equally spaced time points following the cosine curves $cos(2\pi)$ and $cos(4\pi)$, within $[0, 1)$. We use sinusoids of slightly different frequencies at the locations in the same non-zero area to generate higher rank signals. As demonstrated in Table 1 (middle), our method performs well for complex, higher-rank signals. Other regularized regression methods give declining performance for complex signals.

5.1.3 Amount of Spatio-Temporal Correlation. We compare the following three cases: regression covariates with no spatio-temporal correlation, i.e., $\mathbf{X}_i \overset{iid}{\sim} N(0, \mathbf{I} \otimes \mathbf{I})$, small spatio-temporal correlation ($\theta_T = 100, \theta_L = 1$), and large spatio-temporal correlation ($\theta_T = 200, \theta_L = 2$). Table 1 (bottom) shows that when there is no correlation in the data, our method suffers in terms of

Table 1: Linear regression results for varying (top) number of locations, with $n = 100$, $\tau = 200$, $rank(\mathbf{B}_o) = 2$, $SNR = 10$; (middle) rank of signal, with $n = 100$, $\tau = 200$, $L = 100$, $SNR = 10$; and (bottom) spatio-temporal correlation, with $n = 100$, $\tau = 200$, $L = 100$, $rank(\mathbf{B}_o) = 2$, $SNR = 10$.

		Method	Prediction Error	MSE(B)	FPR	TPR
L	25	Loc-Agg	0.4857(0.0208)	0.7932(0.0174)	47.52%(4.22%)	94%(2.09%)
		Lasso	0.6978(0.0307)	1.1896(0.0095)	47.33%(3.10%)	88.50%(3.99%)
		Ridge	0.5789(0.0161)	0.7761(0.0098)	/	/
		Elastic Net	0.5448(0.0218)	0.9871(0.0082)	81.71%(1.31%)	100.00%(0%)
		Tensor Reg.	0.3913(0.0200)	0.8645(0.0165)	/	/
	64	Loc-Agg	0.5290(0.0217)	0.7891(0.0180)	29.29%(2.85%)	92.25%(2.11%)
		Lasso	0.8289(0.0274)	1.1424(0.0064)	23.07%(1.91%)	69.25%(4.24%)
		Ridge	0.7026(0.0193)	0.8358(0.0108)	/	/
		Elastic Net	0.7340(0.0252)	1.0441(0.0099)	51.86%(1.00%)	96.75%(0.78%)
		Tensor Reg.	0.5554(0.0270)	0.9358(0.0179)	/	/
	100	Loc-Agg	0.5209(0.0275)	0.8071(0.0177)	19.98%(3.67%)	88.50%(2.47%)
		Lasso	0.8989(0.0300)	1.1151(0.0087)	12.76%(1.24%)	56.25%(4.52%)
		Ridge	0.7768(0.0212)	0.8995(0.0115)	/	/
		Elastic Net	0.7807(0.0253)	1.0702(0.0092)	38.52%(0.78%)	92.50%(1.43%)
		Tensor Reg.	0.6952(0.0254)	1.0165(0.0107)	/	/
	144	Loc-Agg	0.5889(0.0284)	0.8409(0.0182)	16.56%(2.78%)	89.00%(2.52%)
		Lasso	0.9438(0.0312)	1.1282(0.0089)	9.01%(0.99%)	49.25%(4.19%)
		Ridge	0.8597(0.0264)	0.9396(0.0108)	/	/
		Elastic Net	0.8461(0.0293)	1.0939(0.0088)	30.97%(0.60%)	87.25%(1.62%)
		Tensor Reg.	0.8193(0.0296)	1.0807(0.0123)	/	/
Rank($\mathbf{B}_o$)	2	Loc-Agg	0.5209(0.0275)	0.8071(0.0177)	19.98%(3.67%)	88.50%(2.47%)
		Lasso	0.8989(0.0300)	1.1151(0.0087)	12.76%(1.24%)	56.25%(4.52%)
		Ridge	0.7768(0.0212)	0.8995(0.0115)	/	/
		Elastic Net	0.7807(0.0253)	1.0702(0.0092)	38.52%(0.78%)	92.50%(1.43%)
		Tensor Reg.	0.6952(0.0254)	1.0165(0.0107)	/	/
	4	Loc-Agg	0.5682(0.0285)	0.8424(0.0197)	15.87%(2.48%)	80.00%(3.09%)
		Lasso	0.9549(0.0270)	1.0952(0.0081)	11.22%(1.35%)	46.75%(4.64%)
		Ridge	0.8009(0.0241)	0.9041(0.0139)	/	/
		Elastic Net	0.8306(0.0277)	1.0465(0.0101)	38.74%(0.84%)	91.50%(1.49%)
		Tensor Reg.	0.7073(0.0277)	1.0104(0.0141)	/	/
	8	Loc-Agg	0.6011(0.0251)	0.8170(0.0161)	20.17%(2.91%)	86.50%(2.92%)
		Lasso	0.9446(0.0295)	1.1075(0.0086)	10.20%(1.43%)	50.25%(4.76%)
		Ridge	0.8352(0.0227)	0.9018(0.0130)	/	/
		Elastic Net	0.8552(0.0251)	1.0612(0.0098)	37.72%(0.86%)	94.00%(1.30%)
		Tensor Reg.	0.7594(0.0238)	1.0227(0.0134)	/	/
Correlation	None	Loc-Agg	1.0079(0.0211)	1.0844(0.0076)	77.52%(5.52%)	79.00%(5.35%)
		Lasso	1.0420(0.0218)	1.0404(0.0073)	4.17%(1.01%)	4.25%(1.27%)
		Ridge	1.0281(0.0208)	1.1199(0.0100)	/	/
		Elastic Net	1.0549(0.0213)	1.1273(0.0073)	84.28%(0.61%)	86.75%(1.53%)
		Tensor Reg.	1.0016(0.0204)	1.268(0.0123)	/	/
	Small	Loc-Agg	0.7402(0.0240)	0.9052(0.0121)	41.41%(3.93%)	89.00%(2.87%)
		Lasso	1.0008(0.0199)	1.0775(0.0104)	7.89%(1.38%)	30.75%(4.44%)
		Ridge	0.9225(0.0200)	1.0054(0.0104)	/	/
		Elastic Net	0.9419(0.0224)	1.1059(0.0077)	42.96%(0.84%)	84.00%(1.75%)
		Tensor Reg.	0.8925(0.0159)	1.1419(0.0090)	/	/
	Large	Loc-Agg	0.5209(0.0275)	0.8071(0.0177)	19.98%(3.67%)	88.50%(2.47%)
		Lasso	0.8989(0.0300)	1.1151(0.0087)	12.76%(1.24%)	56.25%(4.52%)
		Ridge	0.7768(0.0212)	0.8995(0.0115)	/	/
		Elastic Net	0.7807(0.0253)	1.0702(0.0092)	38.52%(0.78%)	92.50%(1.43%)
		Tensor Reg.	0.6952(0.0254)	1.0165(0.0107)	/	/

feature selection, while Lasso and Elastic Net can detect some non-zero locations. Note that all methods have about the same true and false positive rates when there is no correlation in the data, which indicates that the non-zero locations are chosen at random as expected. When we have highly correlated data, our Local-Aggregate modeling framework has the best prediction error and signal recovery as we directly account for the spatio-temporal structure of the data.

We have also conducted a similar series of classification simulations with differing numbers of location and signal complexities. However, as the message of the results is similar, the classification results are presented in the supplementary material.

6 Case Study: EEG Data

We demonstrate the utility of our methods for modeling multi-subject neuroimaging data through a case study on Electroencephalography (EEG). We use a well-studied data set from http://archive.ics.uci.edu/ml/datasets/EEG+Database. The data set consists of 122 subjects with 77 alcoholics and 45 controls. We study the single-stimulus EEG recordings averaged over 120 trials while subjects were shown an image. The resulting data set consists of 122 subjects by 57 channels (that have known coordinates out of the 64 channels in the data) by 256 time points forming our predictor tensor $\boldsymbol{\mathcal{X}}$. The objective is to use this multi-subject EEG data to predict the binary response, $\mathbf{y}$, indicating the subject's alcoholic status. We apply our logistic Local-Aggregate model to this data as specified in (2.2). For our local penalties, we employ the group lasso penalty to induce sparsity in the brain locations and a roughness penalty on the times series to enforce temporal smoothness. To construct the smoothing matrix of our aggregating penalty, $\mathbf{G}$, we use spherical distances between the locations of the electrodes on the scalp. Specifically, if $D_{l,l'}$ is the polar distance between node l and l', then we define the spatial kernel smoothing weights as $w_{l,l'} = \exp^{-D^2_{l,l'}/\theta}$ (here, we take $\theta = .1$), and $\mathbf{G} = deg(\mathbf{W}) - \mathbf{W}$, where $deg(\mathbf{W})$ is the degree matrix of $\mathbf{W}$ with $deg(\mathbf{W})_{l,l} = \sum_{l'=1}^{L} w_{l,l'}, deg(\mathbf{W})_{l \neq l'} = 0$. We select the λ_{loc} locally via cross-validation, and compute the Local-Aggregate algorithm path via CV to select the optimal overall model.

We compare the prediction accuracy of our method to that of other competing classification techniques using 5-fold cross-validation and repeat this procedure many times. The average misclassification rates along with standard errors are given in Table 3. We compare our Local-Aggregate model using logistic regression to standard classification methods such as the linear SVM, logistic lasso, and logistic elastic net, as well as the tensor regression method (Zhou and

Li, 2014). Cross-validation was used to select all tuning parameters for all methods. Results reveal that our method performs best in terms of classification accuracy. We expect this better prediction accuracy of our method since it is in fact an ensemble of local models, and fully exploits the spatio-temporal structure of the EEG data.

One major motivation of our approach is to achieve more scientifically interpretable results. To this end, we investigate the coefficients, $\mathbf{B}$, estimated via our procedure as well as that of competitors in Figure 2. Note that $\mathbf{B}$ is a $\tau \times L$ matrix of times series by locations. Hence, we plot all 57 times series (rows of $\mathbf{B}$) in the top and right portions of Figure 2 along with five representative scalp maps (columns of $\mathbf{B}$). Notice that due to the ultra high-dimensionality of the data, the elastic net (top right) estimates overly sparse coefficients; on the other hand, the low-rank constraint of tensor regression (bottom right) results in highly variable coefficients. Our method, however, results in interpretable coefficients, as also seen by the time series arrayed on the scalp map (bottom left). Several studies in addiction literature have found both latency delays and amplitude reduction in visually-evoked and event related potentials among alcoholics compared with controls (Holguin et al., 1999; Porjesz and Begleiter, 2003). Our method finds that alcoholics exhibited lower electrical activity compared to controls in the posterior right and left hemispheres, where the visual cortex lies. Moreover, we can see delays in the activation pattern in the visual cortex among alcoholics compared with controls. Both our findings are consistent with the established literature on the effects of alcoholism on the brain activities after a visual stimulus.

Adding regularization terms to enforce smoothness over the brain locations and time series results in less coefficient variability, and thus more interpretable results. Moreover, our algorithm path selects a better model with the proper amount of smoothing over brain locations. Also notice, that by encouraging sparsity through a group-lasso penalty, our method estimates zeros for several portions of the anterior brain indicating that these areas do not exhibit any difference in electrical activity between alcoholics and controls.

Overall, this case study has demonstrated the strengths of our Local-Aggregate modeling framework, namely improved predictive accuracy and scientifically more interpretable results.

7 Discussion

We have proposed a novel Local-Aggregate modeling framework that translates an ultra-high-dimensional tensor problem into an ensemble of local matrix problems of lower dimensions via

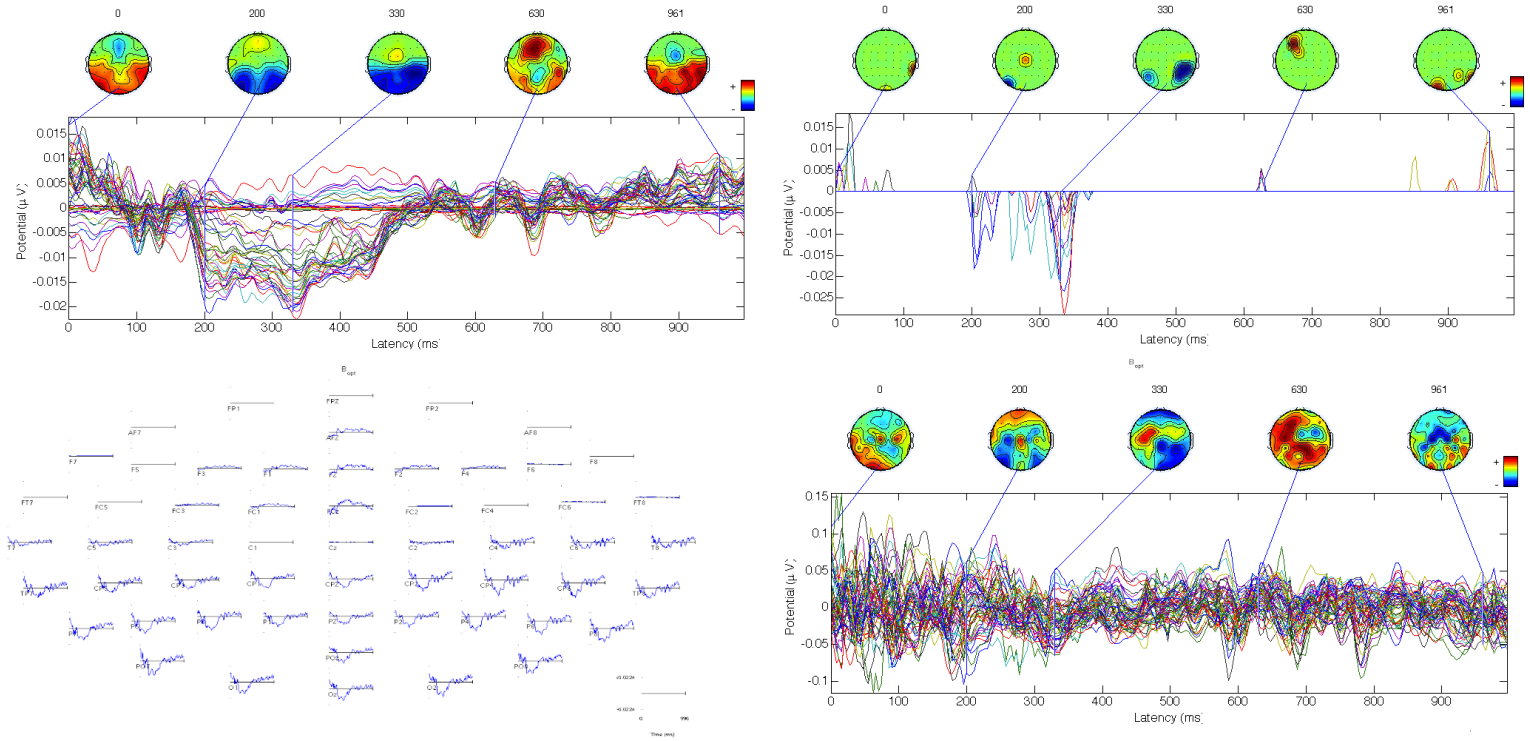

Figure 2: Case study results for the EEG alcoholism study. Left Panel: Coefficients of our Local-Aggregate Model visualized by time series for all locations with four scalp maps shown (top) and with the time series for each location arrayed on a scalp map (bottom). Equivalent time series with the scalp maps are shown for the elastic net (top right) and tensor regression (bottom right). Our method yields scientifically interpretable results, showing that bi-lateral activation in the posterior brain is indicative of alcoholism

Table 2: Prediction Misclassification Error for EEG data

Method	Prediction Error
Loc-Agg	21.65%(0.78%)
Lasso	26.90%(1.19%)
SVM	23.93%(1.11%)
Elastic Net	26.03%(1.06%)
Tensor Reg.	29.32%(1.81%)

regularization. Our modeling framework directly accounts for the spatio-temporal structure of the data through a flexible choice of local penalties and an aggregating penalty that incorporates external spatial information into the modeling framework. From a computational perspective, our highly parallelizable Local-Aggregate ADMM algorithm allows for both distributed memory and data storage, as well as a novel model selection strategy, the Local-Aggregate algorithm path, which chooses the correct model from one of the algorithm iterates, thus greatly reduces the computational burden of model selection. Overall, our Local-Aggregate modeling framework and algorithmic strategy allow us to fit predictive models for large-scale multi-subject neuroimaging data in a fully distributed manner with improved prediction accuracy and scientifically more interpretable coefficients.

While we have presented an EEG application, our methods are ideally suited to predictive models for multi-subject fMRI data. With fMRI data, there are many more locations with often hundreds of thousands of voxels. Based on our experimental results, the larger the number of locations, the better our prediction results. Applications to fMRI data are thus a promising area of future research. Although we are primarily motivated by neuroimaging data, our method is applicable to other structured big data that is collected and stored in a distributed manner, e.g. climate data (weather station), video surveillance data (images), and online shopping data (merchandise). On the computational side, our algorithm is conducive to massive parallel computing via GPUs. Finally statistically, we have introduced a completely novel approach to model selection via the ADMM algorithm path. While empirically we have shown that these well approximate regularization paths, much more work is needed to theoretically study these paths.

In conclusion, our work on Local-Aggregate modeling framework for modeling tensor-valued data has many implications and has opened new possibilities for research both methodologically

and in application to high-dimensional tensor data and model selection.

Acknowledgement

The authors acknowledge support from NSF DMS 1209017, 1317602, and 1264058, and thank Hadley Wickham and Wotao Yin for helpful discussions, and the anonymous reviewers and AE for many helpful suggestions.

Supplementary Material for "Local-Aggregate Modeling for Big-Data via Distributed Optimization: Applications to Neuroimaging"

Yue Hu*

Department of Statistics, Rice University

*email: yue.hu@rice.edu

and

Genevera I. Allen*

Dobelman Family Junior Chair

Departments of Statistics and Electrical & Computer Engineering, Rice University,

Department of Pediatrics-Neurology, Baylor College of Medicine,

Jan and Dan Duncan Neurological Research Institute, Texas Children's Hospital.

*email: gallen@rice.edu

A Supplemental Figures

A.1 Regression Simulation Example

Figure A.1 illustrates the coefficient matrix **B** from the regression simulation example in the paper. We can see that our Local-Aggregate modeling framework correctly selects the locations and achieves locally smooth temporal coefficients at the same time. On the other hand, lasso over-sparsifies the coefficients both in the spatial and temporal domain, regularized tensor regression cannot perform feature selection in the spatial domain.

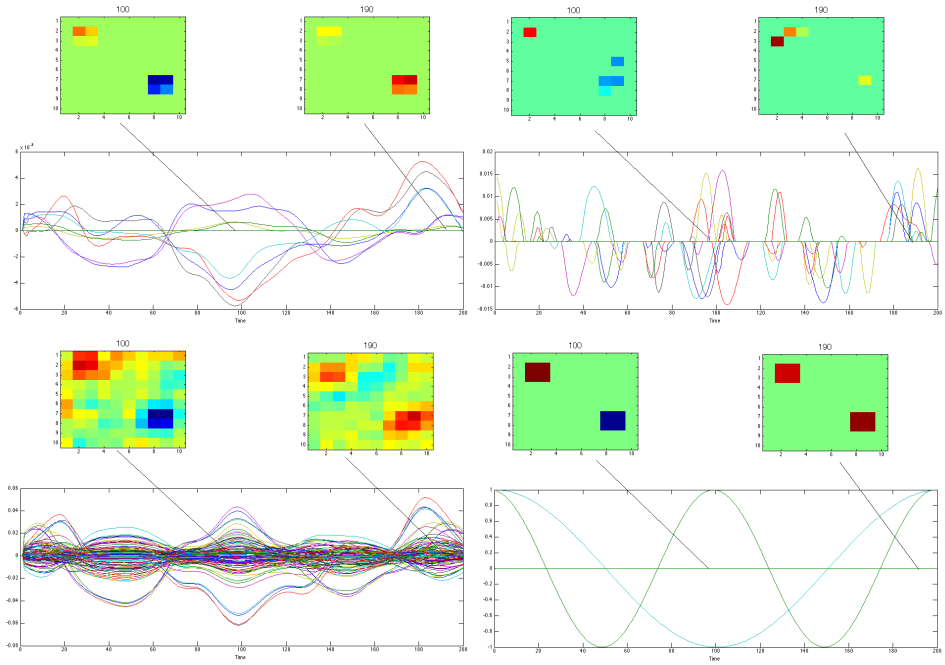

Figure A.1: Comparison of Local-Aggregate method coefficient recovery (top left) with Elastic Net (top right), Tensor Regression (bottom left) and original signal (bottom right). Each subplot shows the time signals at all locations over time with two location signals shown at time point 100 and 190.

A.2 Algorithm Path

We can see in Figure A.2 that the initial, optimal, and final models match exactly for our algorithm path and traditional regularization path in a small simulated example.

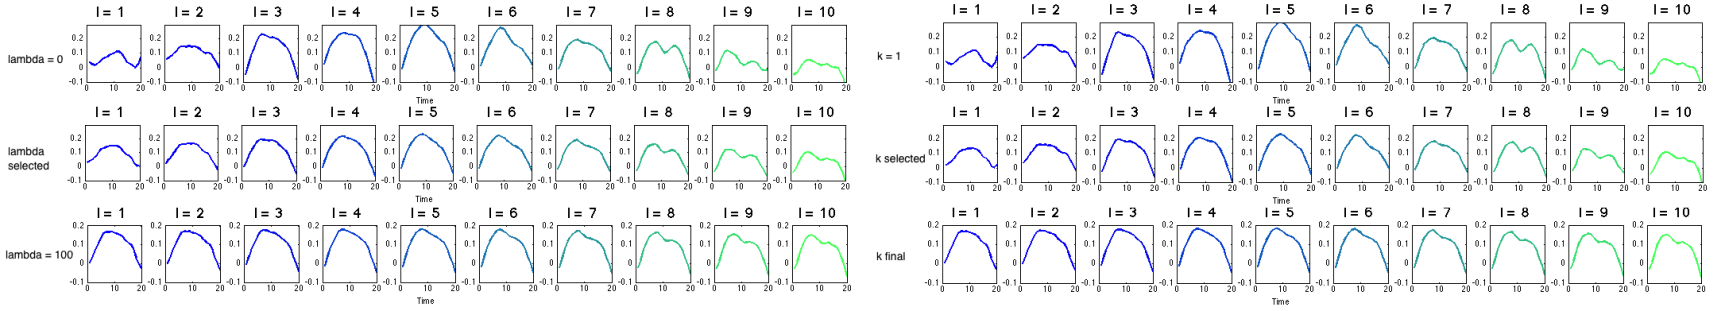

Figure A.2: Comparison of locational coefficients along regularization path (left panels) and algorithm path (right panels) at initial (up), optimal (mid), final (down) stage

B Reparametrizing the Local-Aggregate ADMM framework

We have the following neuroimaging classification problem:

$$\begin{aligned}\underset{\boldsymbol{\alpha},\mathbf{B},\mathbf{Z}}{\text{minimize}} \sum_{l=1}^{L}[\sum_{i=1}^{n}\{-\mathbf{y}_i(\alpha_l+\mathbf{x}_{il}^T\boldsymbol{\beta}_l)+\log(1+e^{(\alpha_l+\mathbf{x}_{il}^T\boldsymbol{\beta}_l)})\}+ \\ \lambda_{loc}^{sm}\boldsymbol{\beta}_l^T\boldsymbol{\Omega}\boldsymbol{\beta}_l+\lambda_{loc}^{sp}\|\boldsymbol{\beta}_l\|_2]+\lambda_{agg}\ \mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T) \\ \text{subject to } \mathbf{B}=\mathbf{Z}\end{aligned} \tag{B.1}$$

However, as the ADMM algorithm only allows for two coupled sets of variables, we need to transform the above problem into the Local-Aggregate ADMM framework with no local intercepts through reparametrization. Let $\widetilde{\mathbf{x}}_{il} = [1;\mathbf{x}_{il}] \in \Re^{(\tau+1)}$, $\widetilde{\boldsymbol{\beta}}_l = [\alpha_l;\boldsymbol{\beta}_l] \in \Re^{(\tau+1)}$, $\mathbf{S} = \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_\tau \end{pmatrix} \in \Re^{(\tau+1)\times(\tau+1)}$, $\widetilde{\boldsymbol{\Omega}}= \begin{pmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Omega} \end{pmatrix} \in \Re^{(\tau+1)\times(\tau+1)}$, we then have an ADMM framework with only 2 variables:

$$\begin{aligned}\underset{\mathbf{B},\mathbf{Z}}{\text{minimize}} \quad & \sum_{l=1}^{L}\big[\sum_{i=1}^{n}\{-\mathbf{y}_i(\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l)+\log(1+e^{\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l})\}+\lambda_{loc}^{sm}\widetilde{\boldsymbol{\beta}}_l^T\widetilde{\boldsymbol{\Omega}}\widetilde{\boldsymbol{\beta}}_l+\lambda_{loc}^{sp}\|\mathbf{S}\,\widetilde{\boldsymbol{\beta}}_l\|_2\big] \\ & +\lambda_{agg}\ \mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T) \\ \text{subject to } & \mathbf{S}\,\widetilde{\mathbf{B}}=\mathbf{Z}\,.\end{aligned} \tag{B.2}$$

Ignoring the $\widetilde{\ }$ sign for notational convenience, the augmented Lagrangian of (B.2) is now given by

$$\begin{aligned}\mathcal{L}_\rho(\mathbf{B},\mathbf{Z},\mathbf{U}) = & \sum_{l=1}^{L}\big[\sum_{i=1}^{n}\{-\mathbf{y}_i(\mathbf{x}_{il}^T\boldsymbol{\beta}_l)+\log(1+e^{\mathbf{x}_{il}^T\boldsymbol{\beta}_l})\}+\lambda_{loc}^{sm}\boldsymbol{\beta}_l^T\boldsymbol{\Omega}\boldsymbol{\beta}_l+\lambda_{loc}^{sp}\|\mathbf{S}\,\boldsymbol{\beta}_l\|_2\big] \\ & +\lambda_{agg}\mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T)+\sum_{l=1}^{L}\frac{\rho}{2}\|\mathbf{S}\,\boldsymbol{\beta}_l-\mathbf{z}_l+\mathbf{u}_l\|_2^2.\end{aligned} \tag{B.3}$$

Lemma B.1. *Minimizing* $\mathcal{L}_\rho(\mathbf{B},\mathbf{Z},\mathbf{U})$ (B.3) *is equivalent to solving the original problem* (B.1)

with local intercept term $\boldsymbol{\alpha}_l$.

Recall that we need to solve the **B**-subproblem, **Z**-subproblem, and then update the dual variable iteratively in order to solve this ADMM optimization framework. We employ the proximal gradient method to solve the local penalized logistic regression as there is a non-smooth local penalty $\| \boldsymbol{\beta}_l \|_2$. Also, the **Z**-subproblem has an analytical solution:

Lemma B.2. *The solution to the* **Z***-subproblem is* $\mathbf{Z}^{k+1} = \rho(\mathbf{S}\,\mathbf{B}^{k+1} + \mathbf{U}^k)\ (2\lambda_{agg} G + \rho\,\mathbf{I}_L)^{-1}$.

Note that the matrix inverse $(2\lambda_{agg} G + \rho\,\mathbf{I}_L)^{-1}$ can be precomputed and since $\mathbf{G}$ is often sparse, the matrix inversion can be done quickly using fast algorithms (see Li et al. (2008)).

C Local Models Vs Multivariate Model

We assume that the generative model for the pair of responses and tensor covariates, $(\mathbf{y}, \boldsymbol{\mathcal{X}})$, follow a matrix generalized linear model (GLM) : $g(\boldsymbol{\mu}) = \alpha + \mathbf{X}_{(1)}^T \text{vec}(\mathbf{B})$, where $\boldsymbol{\mu} = \mathbb{E}(\mathbf{y} \mid \boldsymbol{\mathcal{X}})$ is the conditional mean responses, $\alpha \in \Re$ is the intercept, and $\mathbf{B} \in \Re^{\tau \times L}$ is the coefficient matrix or $\text{vec}(\mathbf{B}) \in \Re^{\tau L}$ is the ultra-high-dimensional coefficient vector. Instead of fitting this model directly, however, we propose to approximate this matrix GLM with a series of local surrogates, $g(\boldsymbol{\mu}_l)$, that we blend together in an ensemble through spatial regularization. The key assumption we make is that conditional on the spatial smoothness of our parameters, the parameters of each local model are such that local models are approximately independent. Beyond this, however, we will see that our local models offer several advantages in terms of mathematical complexity, computational complexity and statistical efficiency.

To understand some of the difference between the matrix GLM and our Local-Aggregate model as well as the advantages of our approach, let us compare methods for the simple case of regression with squared error loss. Here, we will assume we are working with spatio-temporal neuroimaging data and hence employ the same temporal and spatial smoothing penalties. Consider the matrix GLM:

$$\min_{\mathbf{B}} \| \mathbf{y} - \mathbf{X}_{(1)}\, vec(\mathbf{B}) \|^2 + \sum_{l=1}^{L} \lambda_{loc}\, \boldsymbol{\beta}_l^T\, \boldsymbol{\Omega}\, \boldsymbol{\beta}_l + \lambda_{agg} tr(\mathbf{B}\,\mathbf{G}\,\mathbf{B}^T). \tag{C.1}$$

To solve this multivariate model, we rewrite the problem as

$$\min_{\mathbf{B}} \| \mathbf{y} - \mathbf{X}_{(1)} \mathbf{P} \mathbf{B} \mathbf{q} \|^2 + \lambda_{loc} tr(\mathbf{B}^T \mathbf{\Omega} \mathbf{B}) + \lambda_{agg} tr(\mathbf{B} \mathbf{G} \mathbf{B}^T),$$

where $\mathbf{q} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$, $\mathbf{P} = \begin{pmatrix} \mathbf{I} \\ \mathbf{I}_{1,2} \\ \vdots \\ \mathbf{I}_{1,L} \end{pmatrix} \in \Re^{\tau L \times \tau}$, $\mathbf{I}_{1,l}, l = 2, \ldots, L$ are identity matrix with the 1st and l^{th} column switched. Setting the gradient of the multivariate loss function to 0, we have

$$\mathbf{P}^T \mathbf{X}_{(1)}^T (\mathbf{y} - \mathbf{X}_{(1)} \mathbf{P} \mathbf{B}) \mathbf{q}^T + \lambda_{loc} \mathbf{\Omega} \mathbf{B} + \lambda_{agg} \mathbf{B} \mathbf{G} = \mathbf{0}$$

of the form "$\mathbf{A} \mathbf{X} \mathbf{B} + \mathbf{C} \mathbf{X} + \mathbf{X} \mathbf{D} = \mathbf{c}$", which does not have an analytical solution and requires a semidefinite quadratic linear programming solver. On the other hand, our Local-Aggregate model for this problem takes the following form:

$$\min_{\mathbf{B}} \sum_{l=1}^{L} \| \mathbf{y} - \mathbf{X}_l \boldsymbol{\beta}_l \|^2 + \sum_{l=1}^{L} \lambda_{loc} \boldsymbol{\beta}_l^T \mathbf{\Omega} \boldsymbol{\beta}_l + \lambda_{agg} tr(\mathbf{B} \mathbf{G} \mathbf{B}^T) \tag{C.2}$$

Thus we break down the multivariate model into local models that are mathematically much easier to solve and have analytical solutions: $\boldsymbol{\beta}_l = (\mathbf{X}_l^T \mathbf{X}_l + \lambda_{loc} \mathbf{\Omega})^{-1} \mathbf{X}_l \mathbf{y}$. Hence, our Local-Aggregate method is mathematically much easier to solve than the multivariate model.

As discussed in Section 2, we can distribute computations, data storage and memory for our Local-Aggregate model is distributable because our loss is location-seperable: $\ell(\boldsymbol{\mathcal{X}}, \mathbf{B}) = \sum_{l=1}^{L} \| \mathbf{y} - \mathbf{X}_l \boldsymbol{\beta}_l \|_2^2 = \sum_{i=1}^{n} \sum_{l=1}^{L} (y_i - \sum_{t=1}^{\tau} x_{itl} \beta_{tl})^2$. Thus computationally, our model scales well for big data that is collected and stored in a distributed manner. On the other hand, optimization of the multivariate model cannot be parallelized due to the order of summations in the loss function: $\ell(\boldsymbol{\mathcal{X}}, \mathbf{B}) = \| \mathbf{y} - \mathbf{X}_{(1)} \mathrm{vec}(\mathbf{B}) \|_2^2 = \sum_{i=1}^{n} (y_i - \sum_{t=1}^{\tau} \sum_{l=1}^{L} x_{itl} \beta_{tl})^2$. Solving the matrix GLM is then computationally intractable for big data that are too large to load into the memory of one single computer.

The multivariate model vectorizes the $\mathbf{B}$ coefficient matrix which leads to an ultra-high-dimensional problem where $n \ll \tau L$. Because of this, it is not only computationally burdensome, but it is also statistically inefficient to solve. As our Local-Aggregate method breaks down the problem into much more manageable local models of size $n \times \tau$. The sample size complexity

	Method	Prediction Error
L = 25	Loc-Agg	0.3864 (0.0066)
	Matrix GLM	0.6494 (0.0178)
	Local	0.7017 (0.0075)
L = 100	Loc-Agg	1.0408(0.0085)
	Matrix GLM	2.1533 (0.0245)
	Local	1.3002 (0.0098)
L = 144	Loc-Agg	1.2348 (0.0136)
	Matrix GLM	2.2141 (0.0231)
	Local	1.4677 (0.0147)

Table C.1: Comparisons of prediction accuracy for the Local-Aggregate model (Eq C.2), the matrix GLM (Eq C.1), and the Local model (Eq C.3) on the spatial-temporal regression simulation. Here $n = 50, \tau = 50$ and we vary the number of locations L.

n relative to τ versus n relative to τL, is greatly reduces. We then expect higher statistical efficiency in estimating the local models leading to overall improved predictive accuracy with respect to the multivariate model.

We investigate this assertion in a small simulated example. Here, our setup is the same as the that in the regression simulation example presented in Section 5 except we replace the sparse spatial signals with smooth sinusoidal spatial signals, and we use a smaller problem size with $n = 50, \tau = 50$. We compare the different scenarios of $L = 25, 64, 100$, and use the MSE as the criterion for the prediction error. We use the cvx solver for the optimization of the matrix GLM in (C.1).

As indicated in Table C, our Local-Aggregate model has much better prediction error than the matrix GLM in all scenarios and offers substantial improvements when L is larger. We expect that when L is in the tens of thousands as in neuroimaging, our method will yield far better performances as the sample complexity of our local models is constant at τ compared to $L\tau$ for the matrix GLM.

One may wonder that the good performance of the Local-Aggregate method mainly comes from using the separate local loss functions $\sum_{l=1}^{L} \| \mathbf{y} - \mathbf{X}_l \boldsymbol{\beta}_l \|_2^2$ to approximate $\| \mathbf{y} - \mathbf{X}_{(1)} \operatorname{vec}(\mathbf{B}) \|_2^2$, but less from the aggregating penalty, $P_{agg}(\mathbf{B}, \mathbf{G})$, as the Matrix GLM also uses the aggregating penalty. We propose to investigate this question with comparison of our method with the Local model with just the local loss functions and local penalties:

$$\min_{\mathbf{B}} \sum_{l} \| \mathbf{y} - \mathbf{X}_l \boldsymbol{\beta}_l \|_2^2 + \sum_{l=1}^{L} \lambda_{loc} P_l(\boldsymbol{\beta}_l). \tag{C.3}$$

We can see that although the separate loss functions with local penalties alone have better performance compared with the multivariate model when L is large, its performance is not as good as our Local-Aggregate method where both the proposed loss function and aggregating penalties are applied.

In conclusion, compared with the multivariate model, our Local-Aggregate modeling framework is (1) mathematically much easier to solve, (2) computationally more efficient as is conducive to distributed optimization and (3) yields improved statistical efficiency.

D Convergence Accelerators for Local-Aggregate ADMM Algorithm

The stopping criterion of ADMM algorithms typically depends on the norm of the primal and dual residuals (Boyd et al., 2011; Wahlberg et al., 2012; Annergren et al., 2012). We take the stopping criterion to be $\| \mathbf{r} \| < \epsilon^{pr}$ and $\| \mathbf{s} \| < \epsilon^{dual}$ for some $\epsilon^{pr}, \epsilon^{dual} > 0$ as suggested in Boyd et al. (2011). Additionally, the parameter $\rho > 0$ is taken as fixed throughout the algorithm as recommended by Boyd et al. (2011). Under the assumptions of convex GLM loss and penalties and fixed ρ, the convergence of the algorithm is ensured, and the convergence rate is $\mathcal{O}(1/k)$ (Luo, 2012; Deng and Yin, 2012; Goldstein et al., 2012), where k is the iteration number. This convergence speed is still too slow for the purpose of big data modeling, and can be very computationally expensive, for example, it would require at least 1000 iterations even for a rough accuracy of 1e-3. Therefore we need to further speed up the convergence of the algorithm, and cut down the number of iterations to within hundreds.

The convergence of the ADMM can be improved by adapting the parameter ρ to the magnitude of primal and dual residuals across iterations and thereby reducing the dependency on the initial choice of the penalty parameter. Boyd et al. (2011); He et al. (2000) have proposed the following adaptive penalty parameter scheme that tries to keep the norms of primal redisudal ($\mathbf{r} = \mathbf{B} - \mathbf{Z}$) and dual residual ($\mathbf{s}^k = \mathbf{Z}^k - \mathbf{Z}^{k-1}$) within a certain factor of each other as they both converge to zero:

$$\rho^{k+1} = \begin{cases} \tau^{incr}\rho^k & \text{if } \| \mathbf{r}^k \|_2 > \mu \| \mathbf{s}^k \|_2 \\ \rho^k / \tau^{decr} & \text{if } \| \mathbf{s}^k \|_2 > \mu \| \mathbf{r}^k \|_2 \\ \rho^k & \text{otherwise.} \end{cases},$$

where the two multipliers τ^{incr}, τ^{decr} control the acceleration of convergence: the larger τ^{incr}, τ^{decr}

are, the faster the ADMM converges. One common choice may be $\tau^{incr} = \tau^{decr} = 2$ and $\mu = 10$. Other approaches include monotonically increasing or decreasing ρ updates (He et al., 2000), which depend strongly on the starting values of ρ; the self-adaptive updating scheme automatically adjust for the starting value of ρ. However, these schemes do not respect the spatio-temporal structure of the data for our particular problem. Additionally, slow convergence at one time point or location may encumber the overall convergence speed of the entire procedure.

We propose a new technique with self-adaptive penalty parameters that fully respect the complex data structure. We can further speed up the convergence by taking into account of the temporal structure of the data: instead of using one scalar-value, ρ, for all time points and locations, we use a different ρ_t at each time point. The temporal measurements of neuroimaging data exhibit strong spikes in response to stimuli. Hence given the same step size ρ, the convergence rate for these spiked activation periods is slower than at other times because of larger residuals. By forcing the convergence speed to be relatively faster at time points with larger primal residuals, and relatively slower at time points with larger dual residuals, we obtain more balanced convergence speeds over all time points:

$$\rho_t^{k+1} = \begin{cases} \tau^{incr}\rho_t^k & \text{if } \| \mathbf{r}_t^k \|_2 > \mu \| \mathbf{s}_t^k \|_2 \\ \rho_t^k / \tau^{decr} & \text{if } \| \mathbf{s}_t^k \|_2 > \mu \| \mathbf{r}_t^k \|_2 \\ \rho_t^k & \text{otherwise,} \end{cases}$$

where $\| \mathbf{r}_t \|_2^2 = \sum_{l=1}^L \mathbf{r}_{tl}^2$ is the primal residual at time point t, $\| \mathbf{s}_t \|_2^2 = \sum_{l=1}^L \mathbf{s}_{tl}^2$ is the dual residual at time point t. We switch to a fixed ρ updating scheme, i.e., $\tau^{incr} = \tau^{decr} = 1$ after $K = 1000$ iterations.

In order to coordinate convergence speeds at different time points, we propose a new Augmented Lagrangian for the Local-Aggregate modeling framework:

$$\mathcal{L}_{\boldsymbol{\rho}}(\mathbf{B}, \mathbf{Z}, \mathbf{U}) = \sum_{l=1}^{L} \left[r(\mathbf{y}; \mathbf{X}_l \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l) \right] + \lambda_{agg} P_{agg}(\mathbf{Z}, \mathbf{G}) + \sum_{l=1}^{L} \frac{1}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l + \mathbf{u}_l \|_{\boldsymbol{\rho}}^2, \quad \text{(D.1)}$$

where $\| \boldsymbol{\beta}_l - \mathbf{z}_l + \mathbf{u}_l \|_{\boldsymbol{\rho}}^2 = (\boldsymbol{\beta}_l - \mathbf{z}_l + \mathbf{u}_l)^T diag(\boldsymbol{\rho})(\boldsymbol{\beta}_l - \mathbf{z}_l + \mathbf{u}_l)$.

The $\mathbf{Z}$-subproblem, therefore, no longer has a simple least squares solution in the vector $\boldsymbol{\rho}$ case.

Lemma D.1. *The solution of the* $\mathbf{Z}$*-updates of the Local-Aggregate ADMM is equivalent to the*

solution of

$$diag(\boldsymbol{\rho})\,\mathbf{Z} + \mathbf{Z}(2\gamma\,\mathbf{G}) = diag(\boldsymbol{\rho})(\mathbf{B} + \mathbf{U}),$$

which is in the form of the Sylvester equation $\mathbf{A}\,\mathbf{X} + \mathbf{X}\,\mathbf{B} = \mathbf{C}$.

The Sylvester equation does not have an analytical solution when the rank of either $\mathbf{A}$ or $\mathbf{B}$ is greater than 1. Since there is no analytical solution (Sorensen and Zhou, 2003), the MATLAB numerical solver *lyap* is used.

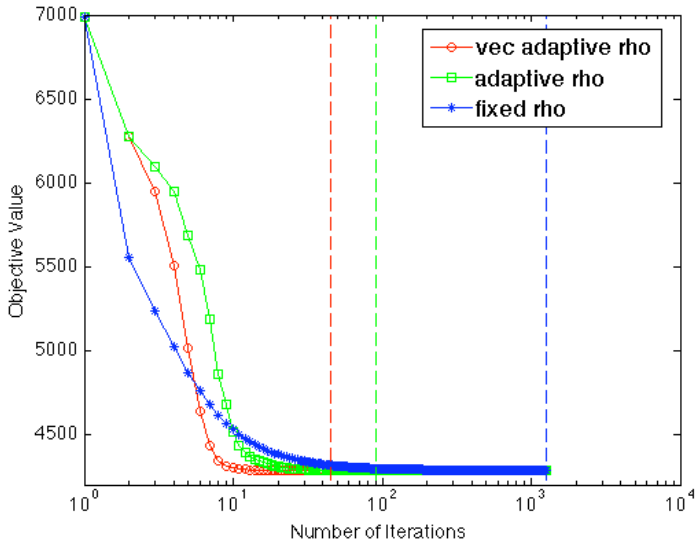

Figure D.3: Comparisons of different updating schemes for $\boldsymbol{\rho}$. The number of iterations till convergence is 45, 91, and 1270 for self-adaptive vectorized $\boldsymbol{\rho}$, self-adaptive scalar ρ and fixed scalar ρ, respectively.

The Local-Aggregate ADMM with the new self-adaptive vectorized $\boldsymbol{\rho}$ penalty parameter is guaranteed to converge, as shown in **Theorem C.2.** We also provide a proof of this theorem in Section E.

As shown in Figure D.3, our self-adaptive vectorized $\boldsymbol{\rho}$ updating scheme yields the fastest monotonic convergence (approximately half of the number of iterations of scalar self-adaptive ρ scheme). The self-adaptive scalar ρ updating scheme is faster in the initial fast-decay stage (the first 5 iterations), and then slows down due to the difference in convergence speeds among time points. The fixed ρ scheme is the slowest amongst all, and depends on the choice of ρ value. Hence, our self-adaptive vectorized $\boldsymbol{\rho}$ updating scheme yields the best convergence performance due to its consideration of the spatio-temporal structure of the data.

To prove convergence of our vector adaptive $\boldsymbol{\rho}$, we write it in general form:

$$\begin{aligned} \underset{\mathbf{B},\mathbf{Z}}{\text{minimize}} \quad & f(\mathbf{B}) + g(\mathbf{Z}) \\ \text{subject to} \quad & \mathbf{B} = \mathbf{Z}, \end{aligned}$$

where

$$f(\mathbf{B}) = \sum_{l=1}^{L} \left[r(\mathbf{y}; \alpha_l + \mathbf{X}_l \, \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l) \right], \qquad g(\mathbf{Z}) = \lambda_{agg} \ \mathrm{tr}(\mathbf{Z} \, \mathbf{G} \, \mathbf{Z}^T).$$

Algorithm 5 then can be put into the following form, letting $\mathbf{H} = diag(\boldsymbol{\rho})$, $f(\mathbf{Z})$ and $g(\mathbf{Z})$ are proper, closed, convex functions.

Algorithm 5 Local-Aggregate Algorithm with Self-Adaptive Vector $\boldsymbol{\rho}$

$$1.\, \mathbf{B}^{k+1} \in \arg\min_{\mathbf{B}} \quad f(\mathbf{B}) + \sum_{l=1}^{L} [(\mathbf{v}_l^k)^T \, \mathbf{H}^k \, \boldsymbol{\beta}_l + \frac{1}{2} \| \, \boldsymbol{\beta}_l - \mathbf{z}_l^k \, \|_{\mathbf{H}^k}^2] \tag{D.2}$$

$$2.\, \mathbf{Z}^{k+1} = \arg\min_{\mathbf{Z}} \quad g(\mathbf{Z}) + \sum_{l=1}^{L} [-(\mathbf{v}_l^k)^T \, \mathbf{H}^k \, \mathbf{z}_l + \frac{1}{2} \| \, \boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l \, \|_{\mathbf{H}^k}^2] \tag{D.3}$$

$$3.\, \mathbf{V}^{k+1} = \mathbf{V}^k + (\mathbf{B}^{k+1} - \mathbf{Z}^{k+1}) \tag{D.4}$$

4. Update $\mathbf{H}^k$

We first make a basic assumption so that the problem has an optimal solution with finite objective value.

Assumption 1. Problem D.2 admits a Lagrangian saddle point, i.e., there exist $\mathbf{B}^* \in \text{dom } f, \mathbf{Z}^* \in \text{dom } g$ and $\mathbf{V}^*$ such that $\mathbf{X}^* = \mathbf{B}^*$ and $f(\mathbf{B}^*) + g(\mathbf{Z}^*) \leq f(\mathbf{B}) + g(\mathbf{Z}) + \sum_{l=1}^{L} \mathbf{v}_l^{*T} (\boldsymbol{\beta}_l - \mathbf{z}_l), \forall \, \mathbf{B} \in \text{dom } f, \forall \, \mathbf{Z} \in \text{dom } g$.

Assumption 2. Problems D.2 and D.3 are solvable.

Assumption 3. The eigenvalues of the semi-positive-definite matrix $\mathbf{H}^k$ are uniformly bounded from below away from zero, and with finitely many exceptions, the eigenvalues of $\mathbf{H}^k - \mathbf{H}^{k+1}$ are nonnegative.

The existence of a saddle point is proved in Gabay and Mercier (1976) and the two subproblems are solvable since f and g in our Local-Aggregate framework are proper convex functions. Our updating scheme for $\mathbf{H}$ is such that after $K = 1000$ iterations, $\mathbf{H}^k - \mathbf{H}^{k+1} = 0$, and we switch to an constant $\mathbf{H}$ updating scheme so that Assumption 3 is always satisfied. Following

from Theorem 2.1. in Kontogiorgis and Meyer (1998), we have the following theorem:

Theorem D.2. *If assumptions 1-3 hold, then for any sequence of iterates* $\{\mathbf{B}^k, \mathbf{Z}^k, \mathbf{V}^k, \mathbf{H}^k\}$ *produced by the alternating direction method :*

(i) $\{\mathbf{B}^k, \mathbf{Z}^k\}$ *converges, and the limit satisfies the constraints of* (D.2)

(ii) $\{f(\mathbf{X}^k), g(\mathbf{Z}^k)\}$ *converges to the optimal value of the objective function for problem* (D.2).

(iii) $\{\mathbf{H}^k\, \mathbf{V}^k\}$ *converges to an optimal dual multiplier for problem* (D.2).

(iv) Any minimizers of problems of the form (D.3) *and* (D.4) *in which* $\mathbf{H}^k\, \mathbf{V}^k$, $\mathbf{B}^{k+1}$ *and* $\mathbf{Z}^k$ *are fixed at their limit values are optimal for problem* (D.2).

Since Assumptions 1, 2 and 3 are all satisfied, our Local-Aggregate algorithm converges. Although there is currently no theoretical proof of a faster convergence speed for the adaptive penalty parameter, the convergence speed is empirically improved as illustrated in Figure D.3.

E Classification Simulation

We will further explore the performance of Local-Aggregate modeling framework for classification problems with local Logistic loss functions. We use the same simulation setup as previously described in the regression simulations in Section 5, except we generate the binary response variable $\mathbf{y}_i \sim Bernoulli(\frac{1}{1+e^{-\mathbf{X}_{(1)}\, vec(\mathbf{B})}})$. We compare our method to the linear support vector machine (SVM), the nuclear-norm regularized tensor logistic regression method (Zhou et al., 2013), the logistic Lasso, and logistic Elastic Net. We will also explore the effect of the number of locations and the rank of the true signal.

E.1 Number of Locations. We employ the same simulation setup with different numbers of locations: $L = 25, 100, 144$. Table E (top) suggests that as the number of locations increases, the prediction accuracy of Local-Aggregate method remains approximately the same while other methods' performance deteriorates. The Local-Aggregate method achieves the best signal recovery as well. The Lasso method over-sparsifies the coefficients due the high dimensionally of the data. We can see that the results are less sensitive to the influence of the number of locations

Table E.2: Classification results for varying the (top) Number of Locations, with $n = 200$, $\tau = 200$, and $rank(\mathbf{B}_o) = 2$, and (bottom) the rank of true signal $\mathbf{B}_o$, with $L = 100$.

		Method	Prediction Error	MSE(B)	FPR	TPR
L	25	Loc-Agg	24.26%(0.81%)	0.7973(0.0140)	72.14%(4.25%)	98.50%(0.85%)
		LASSO	33.18%(1.01%)	1.1807(0.0098)	55.14%(3.37%)	89.00%(4.30%)
		Elastic Net	28.22%(0.89%)	0.9723(0.0090)	89.14%(2.77%)	96.00%(2.80%)
		SVM	25.94%(0.68%)	1.6177(0.0105)	/	/
		Tensor Reg.	23.66%(0.68%)	0.9264(0.0186)	/	/
	100	Loc-Agg	23.78%(0.74%)	0.7806(0.0161)	28.72%(3.90%)	97.25%(0.06%)
		LASSO	40.34%(0.80%)	1.1085(0.0095)	15.24%(1.50%)	60.75%(4.78%)
		Elastic Net	36.70%(0.86%)	1.0599(0.0093)	45.28%(3.40%)	84.50%(4.69%)
		SVM	34.02%(0.74%)	1.6780(0.0065)	/	/
		Tensor Reg.	30.54%(0.83%)	0.9861(0.0125)	/	/
	144	Loc-Agg	24.98%(0.90%)	0.8010(0.0148)	21.31%(3.11%)	91.00%(2.42%)
		LASSO	42.82%(0.93%)	1.0798(0.0114)	8.01%(1.22%)	37.25%(4.90%)
		Elastic Net	38.22%(1.05%)	1.0523(0.0076)	40.40%(2.83%)	84.50%(4.55%)
		SVM	35.58%(0.73%)	1.6611(0.0068)	/	/
		Tensor Reg.	34.36%(0.94%)	1.0365(0.0093)	/	/
Rank ($\mathbf{B}_o$)	2	Loc-Agg	23.78%(0.74%)	0.7806(0.0161)	28.72%(3.90%)	97.25%(0.06%)
		LASSO	40.34%(0.80%)	1.1085(0.0095)	15.24%(1.50%)	60.75%(4.78%)
		Elastic Net	36.70%(0.86%)	1.0599(0.0093)	45.28%(3.40%)	84.50%(4.69%)
		SVM	34.02%(0.74%)	1.6780(0.0065)	/	/
		Tensor Reg.	30.54%(0.83%)	0.9861(0.0125)	/	/
	8	Loc-Agg	22.58%(0.79%)	0.7743(0.0178)	19.43%(3.59%)	91.25%(2.15%)
		LASSO	39.00%(0.97%)	1.1043(0.0091)	15.57%(1.53%)	60.25%(4.88%)
		Elastic Net	34.60%(0.95%)	1.0221(0.0099)	48.00%(2.99%)	92.75%(3.07%)
		SVM	32.44%(0.58%)	1.6916(0.0067)	/	/
		Tensor Reg.	28.74%(0.74%)	0.9594(0.0116)	/	/

than the regression simulation results, since the response variable only depends on the sign of the linear predictor $\mathbf{X}_{(1)}\, vec(\mathbf{B})$.

E.2 Signal Complexity. Now we investigate the behavior of our method under different ranks of the signals given the same number of locations and amount of spatio-temporal correlation. Table E (bottom) suggests that the larger rank of the spatial-temporal signal, the better the detection of the non-zero locations; other regularized regression methods are insensitive to the change of rank of $\mathbf{B}_o$.

F Proofs

Theorem 3.1.

Proof. Boyd et al. (2011) and Mota et al. (2011) together gives a proof of the ADMM algorithm. Let

$$f(\mathbf{B}) = \sum_{l=1}^{L} r(\mathbf{y}; \mathbf{X}_l\, \boldsymbol{\beta}_l) + \lambda_{loc} P_{loc}(\boldsymbol{\beta}_l),$$
$$g(\mathbf{Z}) = \lambda_{agg} P_{agg}(\mathbf{Z}, \mathbf{G}),$$

we can rewrite the Local-Aggregate Framework as

$$\underset{\mathbf{B},\mathbf{Z}}{\text{minimize}} \quad f(\mathbf{B}) + g(\mathbf{Z}) \quad \text{subject to} \quad \mathbf{B} = \mathbf{Z}\,. \tag{F.1}$$

The augmented Lagrangian of (F.1) is

$$\mathcal{L}_\rho(\mathbf{B}, \mathbf{Z}, \mathbf{V}) = f(\mathbf{B}) + g(\mathbf{Z}) + + \sum_{l=1}^{L} [\mathbf{v}_l^T(\boldsymbol{\beta}_l - \mathbf{z}_l) + \frac{\rho}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l \|_2^2].$$

Boyd et al. (2011) shows that if f and g are closed, proper, and convex, and the unaugmented Lagrangian $\mathcal{L}_0$ has a saddle point $(\mathbf{B}^*, \mathbf{Z}^*, \mathbf{V}^*)$, then we have primal residual convergence, i.e., $\mathbf{r}^k = \mathbf{B}^k - \mathbf{Z}^k \to 0$ as $k \to \infty$, dual residual convergence, i.e., $\mathbf{s}^k = \rho(\mathbf{Z}^k - \mathbf{Z}^{k-1}) \to 0$, and also objective convergence, i.e., $p^k \to p^*$, where $p^k = f(\mathbf{B}^k) + g(\mathbf{Z}^k)$. Mota et al. (2011) further shows that the dual variable $\mathbf{V}^k$ converges to the dual optimal point $\mathbf{V}^*$. The only difference from the proof in Boyd et al. (2011) and Mota et al. (2011) is to define a different Lyapunov function of

the algorithm:

$$Q^k = \sum_{l=1}^{L}[(1/\rho)\| \mathbf{v}_l^k - \mathbf{v}_l^* \|_2^2 + \rho\| \mathbf{z}_l^k - \mathbf{z}_l^* \|_2^2].$$

The convergence results of relies on proving the following three key inequalities, and the proof details are the same as in Boyd et al. (2011) and Mota et al. (2011) except for the different Q function.

$$Q^{k+1} \leq Q^k - \sum_{l=1}^{L}[\rho\| \boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1} \|_2^2 + \rho\| \mathbf{z}_l^{k+1} - \mathbf{z}_l^k \|_2^2]. \tag{F.2}$$

$$p^{k+1} - p^* \leq \sum_{l=1}^{L}[-(\mathbf{v}^{k+1})^T(\boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1}) - \rho(\mathbf{z}_l^{k+1} - \mathbf{z}_l^k)^T(\boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1} + \mathbf{z}^{k+1} - \mathbf{z}^*)], \tag{F.3}$$

$$p^* - p^{k+1} \leq \sum_{l=1}^{L}(\mathbf{v}_l^*)^T(\boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l^{k+1}). \tag{F.4}$$

□

Lemma B.1.

Proof. We need to reparametrize the problem (B.1) by letting $\widetilde{\mathbf{x}}_{il} = [1\, \mathbf{x}_{il}] \in \Re^{(\tau+1)}$, $\widetilde{\boldsymbol{\beta}}_l = [\alpha_l; \boldsymbol{\beta}_l] \in \Re^{\tau+1}$, $\mathbf{S} = \begin{pmatrix} 0 & \cdots & 0 \\ \mathbf{0} & \mathbf{I}_\tau \end{pmatrix} \in \Re^{(\tau+1)\times(\tau+1)}$, $\widetilde{\boldsymbol{\Omega}} = \begin{pmatrix} 0 & \cdots & 0 \\ \mathbf{0} & \boldsymbol{\Omega} \end{pmatrix} \in \Re^{(\tau+1)\times(\tau+1)}$, problem (B.1) becomes

$$\begin{aligned} \underset{\widetilde{\mathbf{B}}}{\text{minimize}} \sum_{l=1}^{L}[\sum_{i=1}^{n}\{-\mathbf{y}_i(\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l) + \log(1 + e^{\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l})\} + \lambda_{loc}^{sm}\widetilde{\boldsymbol{\beta}}_l^T\widetilde{\boldsymbol{\Omega}}\widetilde{\boldsymbol{\beta}}_l \\ + \lambda_{loc}^{sp}\| \mathbf{S}\,\widetilde{\boldsymbol{\beta}}_l\|_2] + \lambda_{agg}\ \mathrm{tr}((\mathbf{S}\,\widetilde{\mathbf{B}})\,\mathbf{G}(\mathbf{S}\,\widetilde{\mathbf{B}})^T). \end{aligned} \tag{F.5}$$

Let $\widetilde{\mathbf{Z}} = \mathbf{S}\,\widetilde{\mathbf{B}}$, we can set up the Local-Aggregate ADMM as

$$\begin{aligned} \underset{\widetilde{\mathbf{B}},\widetilde{\mathbf{Z}}}{\text{minimize}} \quad & \sum_{l=1}^{L}[\sum_{i=1}^{n}\{-\mathbf{y}_i(\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l) + \log(1 + e^{\widetilde{\mathbf{x}}_{il}^T\widetilde{\boldsymbol{\beta}}_l})\} + \lambda_{loc}^{sm}\widetilde{\boldsymbol{\beta}}_l^T\widetilde{\boldsymbol{\Omega}}\widetilde{\boldsymbol{\beta}}_l \\ & + \lambda_{loc}^{sp}\| \mathbf{S}\,\widetilde{\boldsymbol{\beta}}_l\|_2] + \lambda_{agg}\mathrm{tr}(\widetilde{\mathbf{Z}}\,\mathbf{G}\,\widetilde{\mathbf{Z}}^T) \\ \text{subject to} \quad & \widetilde{\mathbf{Z}} = \mathbf{S}\,\widetilde{\mathbf{B}}. \end{aligned} \tag{F.6}$$

After rename the variables without the $\widetilde{}$ sign, we have problem (B.2) equivalent to (B.1).

□

Lemma B.2.

Proof.

$$\frac{\partial \mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T) + \sum_{l=1}^{L} \frac{\rho}{2} \| \mathbf{z}_l - \mathbf{S}\,\boldsymbol{\beta}_l^{k+1} - \mathbf{u}_l^k \|_2^2}{\partial\, \mathbf{Z}} = 2\lambda_{agg}\,\mathbf{Z}\,\mathbf{G} + \rho(\mathbf{Z} - \mathbf{S}\,\mathbf{B}^{k+1} - \mathbf{U}^k) = 0$$
$$\mathbf{Z}(2\lambda_{agg}\,\mathbf{G} + \rho\,\mathbf{I}_L) = \rho(\mathbf{S}\,\mathbf{B}^{k+1} + \mathbf{U}^k).$$

□

Lemma C.1.

Proof.

$$\frac{\partial(\gamma\ \mathrm{tr}(\mathbf{Z}\,\mathbf{G}\,\mathbf{Z}^T) + \frac{1}{2}\sum_{l=1}^{L} \| \mathbf{z}_l - \boldsymbol{\beta}_l^{k+1} - \mathbf{u}_l^k \|_{\boldsymbol{\rho}}^2)}{\partial\, \mathbf{Z}} = 2\gamma\,\mathbf{Z}\,\mathbf{G} + diag(\boldsymbol{\rho})(\mathbf{Z} - \mathbf{B} - \mathbf{U}),$$

Setting the gradient to 0 and rearranging terms, we obtain The **Z**-subproblem is thus reduced to solving the Sylvester equation, where $\mathbf{A} = diag(\boldsymbol{\rho})$, $\mathbf{B} = 2\gamma\,\mathbf{G}$, and $\mathbf{C} = diag(\boldsymbol{\rho})(\mathbf{B} + \mathbf{U})$. □

Theorom C.2.

Proof. We will sketch the outline the proof as a collection of lemmas adapted from Kontogiorgis and Meyer (1998), go over the major results and skip the details of the proofs.

We begin the proof by showing in **Lemma E.1** that the iterates $\{(\mathbf{B}^k, \mathbf{Z}^k, \mathbf{V}^k)\}$ are bounded. Then we establish that $\{f(\mathbf{B}^k) + g(\mathbf{Z}^k)\}$ converges to the optimal objective value of the problem (F.1) in **Lemma E.2**. We then show **Lemma E.3** that $\{\mathbf{B}^k\}$, $\{\mathbf{Z}^k\}$, and $\{\mathbf{H}^k\,\mathbf{u}^k\}$ converge, and that the limit of $\{\mathbf{H}^k\,\mathbf{u}^k\}$ is an optimal dual for (F.1). Finally, **Lemma E.4** shows how to obtain a primal solution for (F.1) by solving two minimization subproblems, using the limits of $\{\mathbf{B}^k\}$, $\{\mathbf{Z}^k\}$, and $\{\mathbf{H}^k\,\mathbf{u}^k\}$ as fixed terms in the objective.

Lemma F.1. *Let the assumptions of Theorem C.2 hold, then* $\{\mathbf{B}^k\}$, $\{\mathbf{Z}^k\}$ *and* $\{\mathbf{V}^k\}$ *are bounded.*

Lemma F.2. *Let the assumptions of Theorem C.2 hold, then*
(i) $\{\mathbf{V}^{k+1} - \mathbf{V}^k \to 0\}$, $\{\mathbf{Z}^{k+1} - \mathbf{Z}^k\} \to 0$, $\{\mathbf{B}^k - \mathbf{Z}^k\} \to 0$.
(ii) $\{f(\mathbf{B}^k) + g(\mathbf{Z}^k)\}$ *converges to the optimal objective value for problem* (F.1)

Lemma F.3. *Let the assumptions of Theorem C.2 hold, then*
(i) $\{\mathbf{B}^k\}$, $\{\mathbf{Z}^k\}$ *and* $\{\mathbf{V}^k\}$ *converge.*
(ii) The limit of $\{\mathbf{H}^k\,\mathbf{V}^k\}$ *is an optimal dual for problem* (F.1).

Lemma F.4. *Let* $\{\mathbf{Z}^k, \mathbf{H}^k\,\mathbf{V}^k\}$ *converge to* $(\bar{\mathbf{z}}, \bar{\mathbf{v}})$. *Let* $\widetilde{\mathbf{B}}$ *solve*

$$\underset{\mathbf{B}}{minimize} \quad f(\mathbf{B}) + \sum_{l=1}^{L} \bar{\mathbf{v}_l}^T\, \boldsymbol{\beta}_l + \phi_1(\mathbf{B} - \bar{\mathbf{z}}),$$

and let $\widetilde{\mathbf{Z}}$ *solve*

$$\underset{\mathbf{Z}}{minimize} \quad g(\mathbf{Z}) - \sum_{l=1}^{L} \bar{\mathbf{v}_l}^T\, \mathbf{z}_l + \phi_2(\bar{\mathbf{z}} - \mathbf{Z}),$$

in which ϕ_1 *and* ϕ_2 *are continuous positive definite functions. Then* $(\bar{\mathbf{B}}, \bar{\mathbf{Z}})$ *solves* (F.1).

Part (i) of the theorem can be proven by combining part (i) of **Lemma E.2** and part (i) of **Lemma E.3**. Part (ii) of the theorem is proven in part (ii) of **Lemma E.2** and part (iii) is proven in part (ii) of **Lemma E.3**. Part (iv) is a special case of **Lemma E.4**.

The proof of **Theorem C.2.** follows from Kontogiorgis and Meyer (1998) with some minor modification:
- the Augmented Lagrangian of the problem (F.1) :

$$\mathcal{L}(\mathbf{B}, \mathbf{Z}, \mathbf{U}) = f(\mathbf{B}) + g(\mathbf{Z}) + \sum_{l=1}^{L} [(\mathbf{v}_l)^T\, \mathbf{H}(\boldsymbol{\beta}_l - \mathbf{z}_l) + \frac{1}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l \|^2_{\mathbf{H}}],$$

- when proving **Lemma E.1.** (Lemma 2.5 in Kontogiorgis and Meyer (1998)), we use different trick functions

$$J_1(\mathbf{B}) = f(\mathbf{B}), \quad J_2(\mathbf{B}) = \sum_{l=1}^{L} [(\mathbf{v}_l^k)^T\, \mathbf{H}^k\, \boldsymbol{\beta}_l + \frac{1}{2} \| \boldsymbol{\beta}_l - \mathbf{z}_l^k \|^2_{\mathbf{H}^k}],$$

$$J_1(\mathbf{Z}) = g(\mathbf{Z}), \quad J_2(\mathbf{Z}) = \sum_{l=1}^{L} [-(\mathbf{v}_l^k)^T\, \mathbf{H}^k\, \mathbf{z}_l + \frac{1}{2} \| \boldsymbol{\beta}_l^{k+1} - \mathbf{z}_l \|^2_{\mathbf{H}^k}].$$

□

References

Allen, G. I., L. Grosenick, and J. Taylor (2014). A generalized least-square matrix decomposition. *Journal of the American Statistical Association 109*(505), 145–159.

Annergren, M., A. Hansson, and B. Wahlberg (2012). An admm algorithm for solving l1 regularized mpc. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 4486–4491. IEEE.

Boyd, S., N. Parikh, E. Chu, B. Peleato, and J. Eckstein (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning 3*(1), 1–122.

Calhoun, V. D., J. Liu, and T. Adali (2009). A review of group ica for fmri data and ica for joint inference of imaging, genetic, and erp data. *Neuroimage 45*(1), 163–172.

Chen, C., B. He, Y. Ye, and X. Yuan (2013). The direct extension of admm for multi-block convex minimization problems is not necessarily convergent. *Optimization Online*.

De Martino, F., F. Gentile, F. Esposito, M. Balsi, F. Di Salle, R. Goebel, and E. Formisano (2007). Classification of fmri independent components using ic-fingerprints and support vector machine classifiers. *Neuroimage 34*(1), 177–194.

Deng, W. and W. Yin (2012). On the global and linear convergence of the generalized alternating direction method of multipliers. Technical report, DTIC Document.

Eilers, P. H. and B. D. Marx (1996). Flexible smoothing with b-splines and penalties. *Statistical science*, 89–102.

Friedman, J., T. Hastie, and R. Tibshirani (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software 33(1)*, 1–22.

Friston, K. J., A. P. Holmes, K. J. Worsley, J.-P. Poline, C. D. Frith, and R. S. Frackowiak (1994). Statistical parametric maps in functional imaging: a general linear approach. *Human brain mapping 2*(4), 189–210.

Gabay, D. and B. Mercier (1976). A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications 2*(1), 17–40.

Goldstein, T., B. ODonoghue, and S. Setzer (2012). Fast alternating direction optimization methods. *CAM report*, 12–35.

Grosenick, L., B. Klingenberg, K. Katovich, B. Knutson, and J. E. Taylor (2013). Interpretable whole-brain prediction analysis with graphnet. *NeuroImage 72*, 304–321.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The elements of statistical learning: Data Mining, Inference, and Prediction 2nd edn.* Springer.

He, B., H. Yang, and S. Wang (2000). Alternating direction method with self-adaptive penalty parameters for monotone variational inequalities. *Journal of Optimization Theory and applications 106*(2), 337–356.

Hoerl, A. E. and R. W. Kennard (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics 12*(1), 55–67.

Holguin, S. R., B. Porjesz, D. B. Chorlian, J. Polich, and H. Begleiter (1999). Visual p3a in male alcoholics and controls. *Alcoholism: Clinical and Experimental Research 23*(4), 582–591.

Huang, J. Z., H. Shen, and A. Buja (2009). The analysis of two-way functional data using two-way regularized singular value decompositions. *Journal of the American Statistical Association 104*, 1609–1620.

Hung, H. and C.-C. Wang (2013). Matrix variate logistic regression model with application to eeg data. *Biostatistics 14*(1), 189–202.

Kontogiorgis, S. and R. R. Meyer (1998). A variable-penalty alternating directions method for convex optimization. *Mathematical Programming 83*(1-3), 29–53.

Li, S., S. Ahmed, G. Klimeck, and E. Darve (2008). Computing entries of the inverse of a sparse matrix using the find algorithm. *Journal of Computational Physics 227*(22), 9408–9427.

Luo, Z.-Q. (2012). On the linear convergence of the alternating direction method of multipliers. *arXiv preprint arXiv:1208.3922*.

McCullagh, P. (1984). Generalized linear models. *European Journal of Operational Research 16*(3), 285–292.

Mota, J. F., J. M. Xavier, P. M. Aguiar, and M. Püschel (2011). A proof of convergence for the alternating direction method of multipliers applied to polyhedral-constrained functions. *arXiv preprint arXiv:1112.2295*.

Nesterov, Y. et al. (2007). Gradient methods for minimizing composite objective function. *Core Report, available at http://www.ecore.be/DPs/dp 1191313936.pdf*.

Pereira, F., T. Mitchell, and M. Botvinick (2009). Machine learning classifiers and fmri: a tutorial overview. *Neuroimage 45*(1), 199–209.

Porjesz, B. and H. Begleiter (2003). Alcoholism and human electrophysiology. *Alcohol research and health 27*(2), 153–160.

Ramsay, J. and B. Silverman (2005). *Functional data analysis, 2nd edition*. Springer, New York.

Sorensen, D. C. and Y. Zhou (2003). Direct methods for matrix sylvester and lyapunov equations. *Journal of Applied Mathematics 2003*(6), 277–303.

Tian, T. S., J. Z. Huang, H. Shen, Z. Li, et al. (2012). A two-way regularization method for meg source reconstruction. *The Annals of Applied Statistics 6*(3), 1021–1046.

Wahlberg, B., S. Boyd, M. Annergren, and Y. Wang (2012). An admm algorithm for a class of total variation regularized estimation problems. *arXiv preprint arXiv:1203.1828*.

Yuan, M. and Y. Lin (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68*(1), 49–67.

Zhou, H. and L. Li (2014). Regularized matrix regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76*(2), 463–483.

Zhou, H., L. Li, and H. Zhu (2013). Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association 108*(502), 540–552.